

RL-TR-97-193
Final Technical Report
October 1997



INTEGRATED OPTICAL PROCESSOR

Northrup Grumman Corporation

D.K. Sloper

DTIC QUALITY INSPECTED 2

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.


19980317 060

**Rome Laboratory
Air Force Materiel Command
Rome, New York**

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RL-TR-97-193 has been reviewed and is approved for publication.

APPROVED: 
PAUL L. REPAK
Project Engineer

FOR THE DIRECTOR: 
DONALD W. HANSON, Director
Surveillance & Photonics Directorate

If your address has changed or if you wish to be removed from the Rome Laboratory mailing list, or if the addressee is no longer employed by your organization, please notify RL/OCPC, 25 Electronic Pky, Rome, NY 13441-4515. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE Oct 97		3. REPORT TYPE AND DATES COVERED Final Sep 95 - Apr. 97
4. TITLE AND SUBTITLE INTEGRATED OPTICAL PROCESSOR			5. FUNDING NUMBERS C - F30602-95-C-0023 PE - 632863 PR - 2863 TA - 92 WU - 67	
6. AUTHOR(S) D. K. Sloper				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Northrup Grumman Corporation Electronic Sensors and Systems Division Linthicum, MD 21090			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Rome Laboratory/OCPC 25 Electronic Parkway Rome NY 13441-4515			10. SPONSORING / MONITORING AGENCY REPORT NUMBER RL-TR-97-193	
11. SUPPLEMENTARY NOTES Rome Laboratory Project Engineer: Paul L. Repak, RL/OCPC, (315)330-3146				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The Integrated Optical Processor (IOP) program was designed to investigate the applicability of optical and hybrid opto-electronic approaches to near term surveillance platform processing requirements. The goal of the IOP is the exploitation of the inherent advantages in both optical and electronic technologies to deliver increased throughput with reduced size, weight, and power as compared to an all-electronic implementation. As a target for technology insertion, the Enhanced Objective Airborne Warning and Control System (EOA-AWACS), a major upgrade to the AWACS platform, was identified as a likely candidate to benefit significantly from opto-electronic technology application. Since Space-Time Adaptive Processing (STAP) will play a key role in improving the performance of AWACS, the Lincoln Labs Mountain Top STAP algorithm program was used for the IOP processor analysis model. Using a commercial off the shelf (COTS) approach, a system was designed to be amenable to use as an airborne real-time STAP processor. Results are tabulated and documented which conclude COTS opto-electronic technology, which if implemented into an IOP design, would yield more than an order of magnitude improvement over an all electronic configuration in each volume/weight/power figure of merit category benchmarked.				
14. SUBJECT TERMS optical processing, optical computer architecture, optical switching, micro optical circuits			15. NUMBER OF PAGES 88	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

1. Integrated Optical Processor Final Report

The Integrated Optical Processor (IOP) program was chartered to design a hybrid opto-electronic processor meeting the needs of next generation C³I and surveillance systems. The goal of this design is the exploitation of the inherent advantages in each of the two technologies to deliver increased throughput with reduced size, weight, and power as compared to an all-electronic implementation.

Two teams, each led separately by Northrop-Grumman Corporation, were tasked with assessing the emerging needs of the AWACS and J-STARS systems, to identify how and where the insertion of opto-electronic technologies can best contribute to the improved performance and continued viability of these two systems well into the 21st century.

This report documents the work performed by Northrop-Grumman's Electronic Sensors and Systems Division of Baltimore, Maryland in support of the IOP program from October 1995 through April, 1997. In addition, the following subcontractors contributed significantly to the IOP study:

Honeywell's Systems and Research Center in Bloomington, Minnesota provided expertise in the application of optical interconnects within systems, performed optical technology assessments and designed the opto-electronic crossbar switch which is an essential element in the IOP design.

Mercury Computer Corporation, of North Chelmsford, Massachusetts provided expertise in the area of real-time multi-computing systems, and contributed substantially to our selection of an appropriate architecture for the IOP. This architecture was found to be heavily influenced by the use of optical interconnects.

Syracuse Research Corporation, of Syracuse, New York provided expertise in Residue Number System (RNS) processing techniques and algorithms, and designed a set of RNS-based hardware accelerator modules, which are an essential element in our IOP design.

TABLE OF CONTENTS

1. INTRODUCTION:	1
1.1 OBJECTIVE	1
1.2 STUDY METHODOLOGY	1
1.3 RESULTS OF THE IOP STUDY.....	3
2. APPLICATION OF THE IOP	5
3. PROCESSING REQUIREMENTS AND ALGORITHMS	8
3.1 IOP PROCESSING REQUIREMENTS:.....	8
3.2 STAP ALGORITHMS:	9
3.3 DEFINING IOP STAP SYSTEM PARAMETERS:	10
3.4 STAP PERFORMANCE GOALS:.....	11
4. COTS-BASED STAP PROCESSING.....	13
4.1 PERFORMANCE OF COTS-BASED STAP ARCHITECTURES:.....	13
4.2 IDENTIFYING LIMITATIONS IN COTS-BASED STAP PROCESSORS:	14
4.3 QUANTIFYING THE IMPACT OF DATA COMMUNICATIONS IN COTS-BASED STAP:.....	16
4.4 LOCATING THE COTS DATA COMMUNICATIONS BOTTLENECK:.....	17
4.5 SOLUTIONS TO THE DATA COMMUNICATIONS BOTTLENECK:.....	18
4.5.1 Resource Concentration (Hardware Acceleration):	19
4.5.2 Increasing Architectural “Robustness” (Topology):.....	19
4.5.3 Increased Link Bandwidth	20
4.6 THE NEED FOR OPTICS.....	21
4.6.1 STAP Accelerator:.....	21
4.6.2 System Topology:	21

4.6.3 Passive Backplanes:	21
4.6.4 Sensor-to-Processor Interconnects:	22
4.7 SUMMARY	22
5. IOP ARCHITECTURE.....	23
5.1 SELECTION OF THE MERCURY RACE AS A BASELINE STAP ARCHITECTURE	23
5.2 ARCHITECTURAL DECISIONS DRIVEN BY THE INTENDED APPLICATION OF THE IOP	24
5.3 ARCHITECTURAL DECISIONS DRIVEN BY THE OPERATING ENVIRONMENT OF THE IOP	25
5.3.1 Airborne Environment:	25
5.3.2 Software:	25
5.3.3 “Core technology” vs. “Off-The-Shelf”	26
5.3.4 Cooling.....	26
5.4 KEY INNOVATIONS REQUIRED FOR IOP	26
5.5 CONCLUSIONS:	27
6. OPTICAL INTERCONNECT TECHNOLOGY.....	30
6.1 OPTICAL LINK COMPONENTS AND BUILDING BLOCKS	30
6.1.1 Vertical cavity surface emitting lasers (VCSELs)	30
6.1.2 Integrated optical receiver.....	30
6.1.3 Polymer optical waveguides:	31
6.1.4 Waveguide connectors	32
6.1.5 Tx/Rx Packaging, and optical links	33
6.2 OPTICAL DATA LINK	34
6.2.1 Serial vs. Parallel:	34
6.2.2 Fiber vs. Waveguide:	35

6.3 OPTICAL INTERCONNECTS FOR IOP MERCURY SYSTEMS.....	35
6.3.1 The Mercury RACE Architecture	36
6.3.2 The IOP STAP Processor.....	37
6.3.3 The IOP backplane optical interconnect link design:	38
6.3.4 Optical RACE router module.....	39
6.3.4.1 Power considerations	40
6.3.4.2 Cost and other considerations	41
6.3.5 Polymer optical waveguide, and backplane connector	41
6.3.5.1 Link loss budget.....	42
6.3.5.2 Bit error rate (BER)	43
6.3.5.3 Board-to-backplane connector, and its density	43
6.4 IOP OPTICAL BACKPLANE INTERCONNECT RELIABILITY ISSUES	44
6.4.1 VCSELs	44
6.4.2 Integrated Receiver	46
6.4.3 Optical Router Chip	46
6.4.4 Optical polymer waveguides and backplane connectors	47
6.5 SUMMARY.....	47
7. ACCELERATOR NODES	48
7.1 DESIGN TRADES	51
7.1.1 RNS Arithmetic	52
7.1.2 RNS Choleski Algorithm.....	53
7.1.3 RNS STAP Implementation.....	54
7.2 CORNER TURN MEMORY	55
7.3 FFT MODULE.....	56

7.4 WEIGHT MAKER MODULE	57
7.4.1 Matrix Multiplication ASIC.....	59
7.4.2 Core ASIC.....	59
7.4.3 Power Dissipation For Weight Maker Module	61
7.5 BEAMFORMER MODULE	61
8. DESIGN PARAMETERS:	69
8.1 SYSTEM PERFORMANCE:	69
8.2 SUMMARY:.....	71

TABLE OF FIGURES

Figure 1.2-1	IOP Processor Design Methodology -----	1
Figure 1.3-1	Relative Performance of IOP vs. a Conventional COTS-Based Design-----	4
Figure 2.1	Enhanced Objective AWACS System Configuration-----	6
Figure 2.2	Enhanced Objective AWACS Processor Architecture-----	7
Figure 3.1	HOPD STAP Algorithm Flow Diagram-----	10
Figure 4.1-1	Achieved COTS STAP Processing Latency-----	14
Figure 4.2-1	Achieved Speed Up of Various STAP Partitionings-----	15
Figure 4.3-1	Contribution to STAP Execution Time of Constituent Sub-functions--	16
Figure 4.4	A Distributed Corner Turn-----	18
Figure 5.2-1	STAP Processing Characterized-----	24
Figure 5.4-1	IOP Advances to OMNET Optical Backplane Technology-----	27
Figure 5.5-1	IOP Mercury Based STAP Architecture-----	28
Figure 6.1.1	32-Channel VCSEL Array-----	30
Figure 6.1-2	GaAs Foundry Processed Integrated Rx Layout, Chip, and Package--	31
Figure 6.1-3	12-ch. Integrated Rx Design for 2.5Gbps Per Each Channel-----	31
Figure 6.1-4	Eye Diagram at 1.6 Gbps Honeywell VCSEL → Ga As COTS Rx -----	31
Figure 6.1-5	(a) Cross Section Microphotograph (b) Photo of ULTEM/BCB polymer waveguide	

TABLE OF FIGURES (CONTINUED)

	fabricated on Kapton substrate-----	32
Figure 6.1-6	Waveguide MCM To Board Connector With Passive Alignment Feature-	32
Figure 6.1-7	Board-to-Backplane Connect Based On Expanded Beam Approach -----	33
Figure 6.1-8	OETC 32-Channel Parallel Fiber Link Tx Modules-----	33
Figure 6.1-9	(a) Gigabit Tx/Rx Module	
	(b) Three-Channel FODB Tx/Rx MCM	
	(c) HDI Packaged Tx/Rx Array with Connectorized Polymer Waveguides	
	(d) MCM-Ceramic Packaged Multi-channel Tx/Rx With Ploymer	
	Waveguide Interfaces-----	34
Figure 6.2-1	Schematic Block Diagram of a Optical Link-----	34
Figure 6.2-2	Optical Fiber and Waveguides for Parallel Optical Interconnects-----	35
Figure 6.3-1	A 30-Node System in 6U VME Chassis-----	36
Figure 6.3-2	Optical Interconnect for IOP Mercury System-----	37
Figure 6.3-3	A Schematic IOP Processor in a ^U Chassis With Optical Backplane-----	38
Figure 6.3-4	A Schematic Illustration of a Board, and a Router With Optical RACE Ports	39
Figure 6.3-5	A Schematic Illustration of a Optical Router MCM-----	40

TABLE OF FIGURES (CONTINUED)

Figure 6.3-6	A Schematic Illustration of the VCSEL-to Waveguide Interface----	41
Figure 7.0-1	IOP System Block Diagram-----	49
Figure 7.0-2	IOP System Data Rates-----	50
Figure 7.0-3	STAP Data Cube-----	50
Figure 7.1-1	IOP Throughput Requirements-----	51
Figure 7.1.1-1	Factored RNS Multiplier/Adder Schematic-----	53
Figure 7.1.2-1	Choleski RNS Weight Calculation Algorithm-----	54
Figure 7.2-1	Corner Turn Memory Module-----	56
Figure 7.4-1	Weight Maker Module-----	58
Figure 7.4.1-1	Matrix Multiplication ASIC-----	60
Figure 7.4.2-1	CORE ASIC-----	62
Figure 7.4.2-2	RNS Inverse Square Root-----	63
Figure 7.4.2-3	Binary-to-RNS Conversion-----	64
Figure 7.4.2-4	RNS-to-Binary Conversion-----	65
Figure 7.4.2-5	(a) CRNS-to-QRNS Conversion	66
	(b) QRNS-to-CRNS Conversion-----	67
Figure 7.5-1	Beamformer Module-----	68
Figure 8.2-1	Single 6U Chassis IOP STAP Processor-----	71
Figure 8.2-2	Relative Processor Performance-----	72

LISTING OF TABLES

Table 1.3-1	IOP Processor Complexity-----	4
Table 1.3-2	IOP & Mountain Top Processor "Figures of Merit-----	4
Table 3.1	IOP Throughput/Data Rate Requirements-----	8
Table 3.2	Mountain Top/IOP STAP System Parameters-----	11
Table 3.3	Mountain Top/IOP Stap Performance Goals-----	12
Table 6.1	IOP Backplane Optical Link Loss Budget-----	42
Table 6.4.1	Summary of Optical Link Components Reliability-----	46
Table 8.1-1	Mountain Top Processor Estimates-----	70
Table 8.1-2	IOP Processor Estimates-----	70
Table 8.1-3	Performance Comparison-----	70

INTEGRATED OPTICAL PROCESSOR FINAL REPORT

1. INTRODUCTION:

1.1 OBJECTIVE

The performance limits of existing surveillance aircraft (e.g. AWACS, JSTARS) are now being stressed by the emergence of low-observable threats, sophisticated electronic countermeasures, increased target densities, and the complexity of engagement of the modern battlefield.

A number of techniques including fused multi-spectral sensors, adaptive clutter cancellation, and electronic counter-counter measures have been widely identified as means to increase surveillance capabilities against these threats. Processing requirements of many of these schemes, however, remain prohibitive, outpacing the rate of advance of conventional electronics. Hybrid opto-electronic processing systems offer one potential solution to this processing dilemma. The Integrated Optical Processor program was chartered to investigate the applicability of opto-electronics to the surveillance processing problem.

The objective of this effort was to design a hybrid opto-electronic processor which exploits the advantages inherent to each of the two technologies to produce a processor which delivers increased throughput with reduced size, weight, and power, as compared to an all electronic implementation. This report summarizes the design of the IOP, its characteristics and some of the trade-offs made during the design process.

1.2 STUDY METHODOLOGY

Figure 1.2-1 illustrates the design methodology employed for the IOP program.

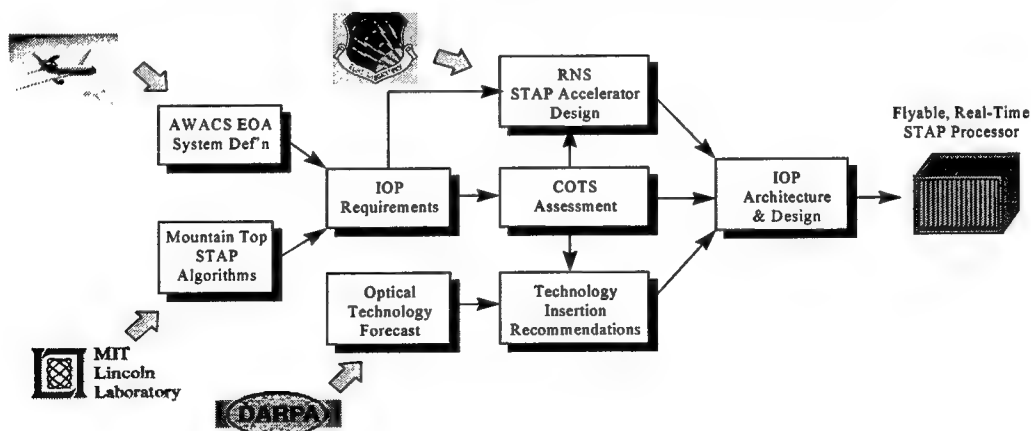


Figure 1.2-1 IOP Processor Design Methodology

To be successful, we believe the design of any future military system must couple the innovations developed under core technology programs with the advantages afforded through the use of "mainstream" COTS technologies. Our IOP design approach is intended to leverage strongly from ongoing efforts in each of these areas.

The Enhanced Objective AWACS (EOA) program, a major upgrade to the AWACS platform was identified as a likely candidate to benefit significantly from opto-electronic technology insertion. Section 2 describes the targeted improvements to the AWACS platform contemplated as part of the EOA program.

Using these proposed improvements to the AWACS system as a guide, we next established processing requirements for the AWACS EOA. Since Space-Time Adaptive Processing (STAP) will play a key role in improving the performance of AWACS against an increasing threat from low-observable targets (most notably low-flying cruise missiles) and electronic countermeasures, we looked to the Lincoln Labs Mountain Top program to provide a set of STAP algorithms for use in our processor requirements analysis. These algorithms and their resulting processing requirements are described in Section 3.

To determine whether a technology insertion is warranted, we next assessed the ability of the next generation of conventional (i.e. all-electronic) COTS-based systems to meet the processing requirements we derived for the AWACS EOA system. This assessment is documented in Section 4.

The conclusions drawn from our COTS technology assessment pointed to data communications as being the key roadblock standing in the way of a COTS-based STAP processor. This led to our search for possible means to solve this data communications bottleneck. Optics became a key element of our solution. Our resulting IOP architecture, and the impact of optical interconnects upon that architecture, are documented in Section 5.

The benefits afforded by a set of emerging optical technologies were found to largely solve our STAP data communications problems, if they could be applied to the IOP design. Our assessment of the "readiness" of these technologies for system insertion, and their incorporation into the design of the IOP are documented in Section 6.

To be viable in an airborne surveillance environment, our STAP solution must not only deliver the necessary processing throughputs, but must also meet stringent size, weight, power, and environmental constraints. A set of STAP-oriented hardware accelerator modules was identified as a means to derive substantial size, weight, power, and cost savings, over a "fully-COTS" based approach. These Residue-Number System (RNS) based accelerators leverage heavily from work performed on earlier phases of the IOP program, and on other efforts targeting use of RNS funded by Rome Labs. Our work in this area proved so successful, that we believe we have opened up the possibility for the insertion of STAP processing into systems such as airships and UAVs. Our STAP accelerator design efforts are described in Section 7.

To measure the success of our IOP design, we again relied upon the work of the Mountain Top program. Using the processor performance metrics predicted for the Mountain Top STAP processor, we quantitatively compared the performance of our IOP processor design to that achieved using a "COTS-only" approach. This IOP performance assessment is documented in Section 8.

1.3 RESULTS OF THE IOP STUDY

The IOP study clearly demonstrates that the insertion of optical interconnect technologies will have a major impact upon the performance of a real-time multiprocessing system.

By analyzing measured results derived as part of the Mountain Top program, our work plainly demonstrates both the need for, and the benefits of, optical insertion into a COTS-based STAP processor. These results have implication well beyond the limited scope of STAP, however, and clearly portray the advantages of optical interconnects in any high performance processor.

Our IOP design approach builds upon COTS processing solutions from Mercury with technology innovations under development in DARPA's OMNET and POINT programs. As part of our IOP design, we propose to advance each of the technologies to its next logical step, combining POINT's polymer waveguides with OMNETS's optoelectronic packaging innovations to produce a flightworthy integrated opto-electronic processing system.

We have shown such a system to be amenable to use as an airborne real-time STAP processor. *Our resulting IOP design provides full STAP capability, with over 150 Giga-ops of sustained throughput in a single 6-U VME chassis.* Our work has thus opened up the possibility for the insertion of STAP processing into systems where it has heretofore proven unfeasible.

Table 1.3-1 lists our estimates for the size, weight and power of the IOP STAP processor. Combining these results with our IOP throughput requirements yields a set of "figures of merit", useful for comparative benchmarking of our design to that of a conventional all-electronic implementation. Table 1.3-2 lists these performance metrics for the IOP, along with similar metrics for the Mountain Top processor.

The relative performance of our approach to that employed on the Mountain Top is portrayed pictorially in Figure 1.3-1. *Our IOP STAP processor is shown to exhibit more than an order of magnitude improvement in each category.*

TABLE 1.3-1 IOP PROCESSOR COMPLEXITY

Size	2.25 ft ³
Weight	125 lbs
Power	760 W
GigaFlops (sustained)	152 Gflops

TABLE 1.3-2. IOP & MOUNTAIN TOP PROCESSOR “FIGURES OF MERIT”

	Mountain Top Phase I Predicted	Mountain Top Phase II Predicted	IOP Predicted	Performance Delta IOP/Mtn Top
Gflops/ft ³	1.53	1.50	67.55	45x
Gflops/lb	0.064	0.053	1.216	23x
Gflops/KW	8.77	3.44	200.0	58x




Figure 1.3-1 Relative Performance of IOP vs. a Conventional COTS-Based Design

2. APPLICATION OF THE IOP

As defined by the IOP Statement of Work, the targeted application of the IOP is next generation airborne surveillance. This could include a variety of surveillance platforms including AWACS, JSTARS, and UAVs such as Tier 2+ and Tier3-. Although all of these systems have the potential for using the IOP, the Northrop Grumman ESSD team is concentrating on AWACS as its primary application (a second Northrop Grumman team is concentrating on J-STARS). However, our decision to base our IOP design upon a COTS architecture will likely make it of interest to these other application as well.

Planned AWACS improvements include automatic target ID, ECCM, Improved Man-Machine Interface and Multisensor Integration including replacement of the existing mission computer. None of these modifications, however, is seen to require a processor of the potential capability of the IOP.

However, there is a new thrust at Northrop-Grumman sponsored by the Electronic Systems Command, termed "Enhanced Objective AWACS" (EOA). This program is being structured so as to add a new UHF radar to the AWACS system to enhance its current detection range against low-observables.

The proposed system configuration will use the existing S-Band antenna and transmitter, augmented will new advanced multichannel receivers and processors. The vision is to include processing for both the S-Band and UHF radars in a single, integrated, high speed processor meeting the size, weight, and power constraints of the existing system. The EOA system configuration is shown in Figure 2.1.

In order to detect low-flying, low radar cross section cruise missiles, this new system will clearly require STAP processing. To deliver the processing throughput needed for STAP while meeting the existing size, weight, and power restrictions will likely require advanced techniques, perhaps including optical interconnects. The proposed EOA processing architecture is shown in Figure 2.2.

We believe our IOP design will meet the needs of the AWACS EOA program, and many additional applications as well. (We are particularly interested in potential UAV markets, due to the exceedingly low size, weight, and power of our design.). In fact, Mercury has indicated that they are mainly interested in the IOP program due to its multi-application potential, thereby helping to meet their next generation processing needs for surveillance and other real-time applications.

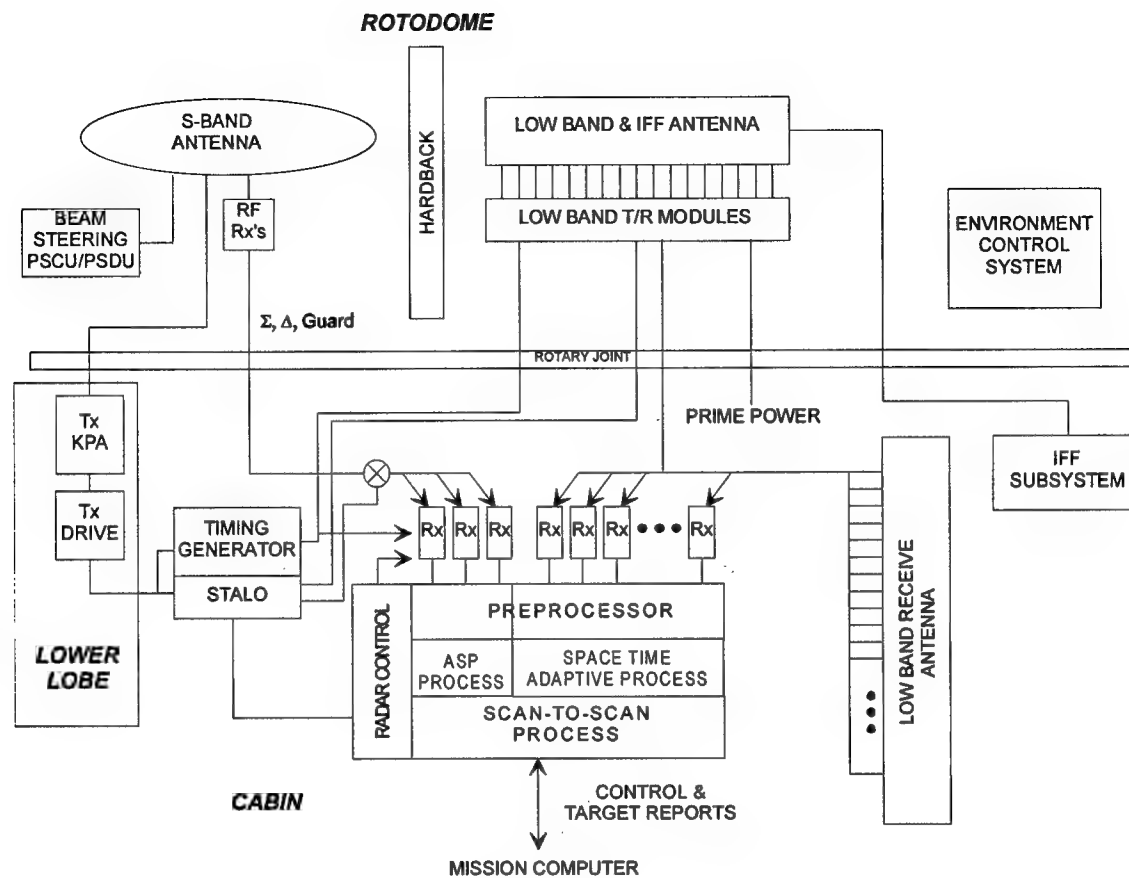


Figure 2.1 Enhanced Objective AWACS System configuration

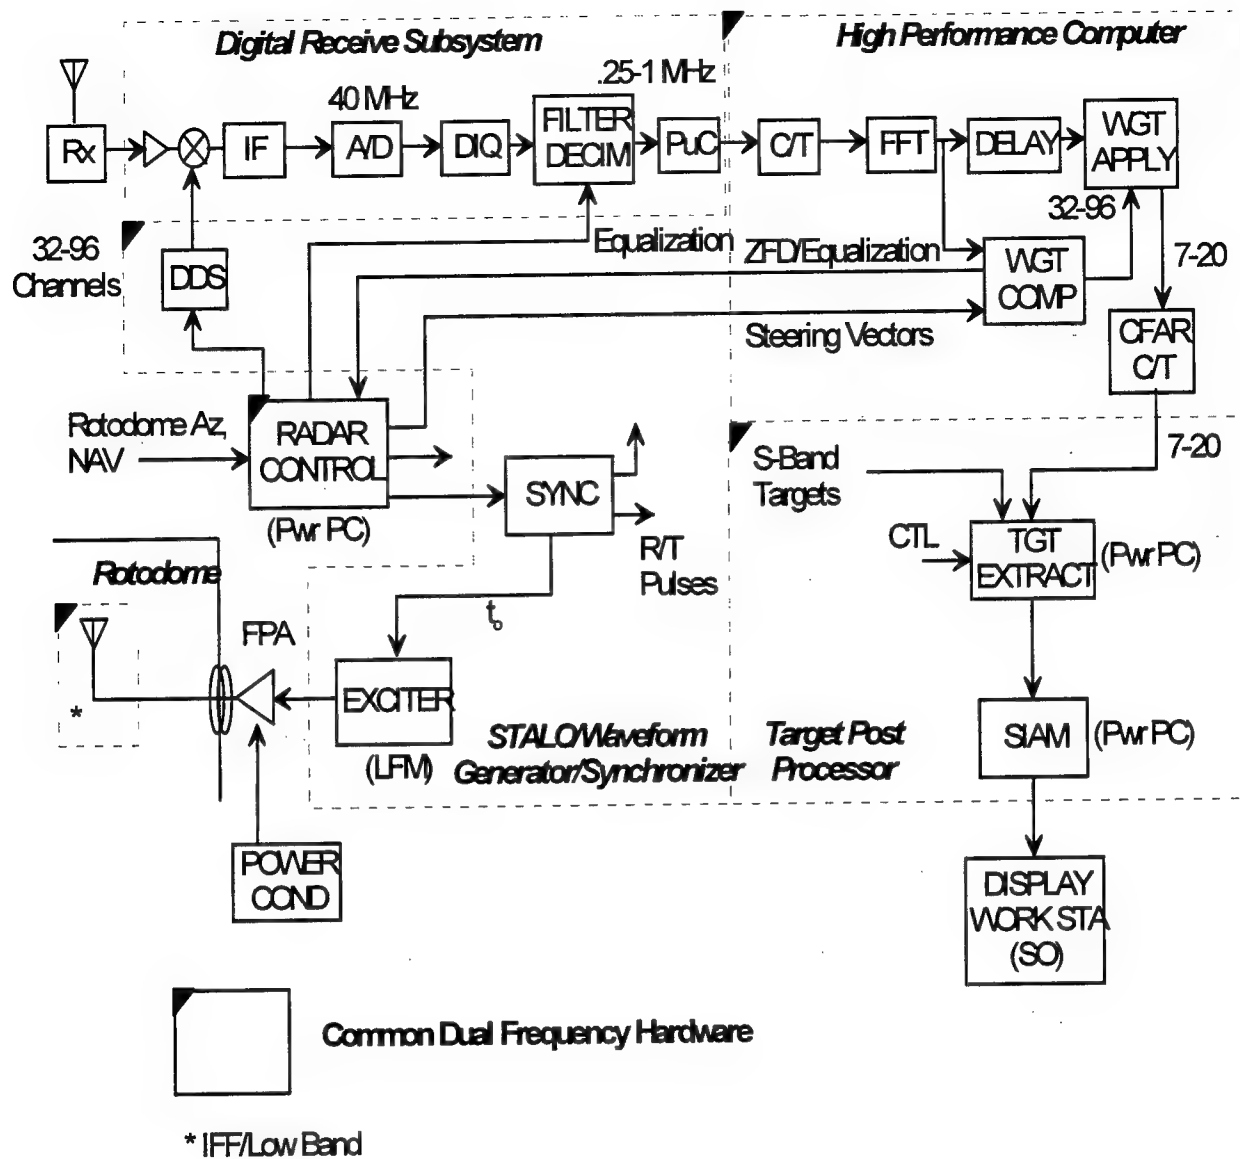


Figure 2.2 Enhanced Objective AWACS Processor Architecture

3. PROCESSING REQUIREMENTS AND ALGORITHMS

Overview:

This section describes the IOP processor throughput and interconnect system bandwidth requirements, the STAP algorithms selected for use in the IOP processor sizing study, and the performance goals established for the IOP STAP processor.

3.1 IOP PROCESSING REQUIREMENTS:

Processor throughput and interconnection system data rate requirements have been estimated for the AWACS EOA application. These are summarized in Table 3.1. Requirements relating to the S-Band radar functions were derived from existing AWACS RSIP requirements, since these functions are likely to be ported intact to the EOA platform. UHF radar requirements were derived using an initial system specification for the EOA's low-band radar and the High Order Post-Doppler (HOPD) STAP algorithm developed as part of the MIT Lincoln Labs Mountain Top program.

TABLE 3.1 IOP THROUGHPUT/DATA RATE REQUIREMENTS

Function	Sub-Function	Throughput (GOPS)	Data Rate (Gbit/sec)
S-Band Radar	Pulse Compression	1.28	0.24
	Clutter Cancellation	1.48	0.28
	FFT	0.4	0.32
	Beamforming	4.8	0.32
	Detect/Centroid	0.58	0.32
	CFAR	0.75	0.32
UHF Radar	Preprocessing	196.6	102.4
	FFT	4.9	5.1
	Weight Generation	140	5.9
	Weight Application	11.8	5.9
	Detect/Centroid	0.1	0.4
	CFAR	0.1	0.4

Table 3.1 shows the throughput and data flow requirements for the EOA platform to be dominated by the STAP and digital preprocessing functions of the UHF radar. While it may be assumed that the remaining requirements will be met by commercial processor technologies available in the EOA development timeframe, the STAP and preprocessor functions bear closer scrutiny, and have been identified as candidates for potential insertion of optical interconnect technologies and Residue Number System (RNS) processing techniques.

3.2 STAP ALGORITHMS:

A significant number of research efforts are currently underway throughout industry and academia in the area of STAP algorithm development. In an effort to benefit from these activities, we have selected a set of three STAP algorithms developed by MIT Lincoln Labs as part of the ongoing DARPA funded Mountain Top program for possible application to the AWACS EOA. (The Mountain Top program's charter includes a mandate to *"assess the suitability of candidate architectures and processors for real-time implementation of STAP algorithms"*. This work is being conducted at the White Sands Missile Test Range in New Mexico, and at the Pacific Missile Range Facility in Hawaii.)

The Mountain Top STAP algorithms were chosen because they represent a *de facto* industry standard, having been extensively benchmarked by a number of independent contractors on a variety of COTS-based platforms including the Embedded Touchstone (Honeywell), Mercury Race (Northrop Grumman), Cray (Cray Research), and Paragon P6 (Intel). These benchmarks thus serve as a useful metric for determining if a potential technology insertion is warranted, and if so, as a tool for quantifying the impact of the insertion. The algorithms under consideration are:

Beam Space PRI Staggered Post-Doppler (BSPD)

The BSPD algorithm performs reduced dimension space-time adaptive nulling. Nulling is achieved by adaptively combining a set of beams using weights computed from a power selected training set. Two overlapped PRI sets are used.

Element Space PRI Staggered Post-Doppler (ESPD)

The ESPD algorithm performs reduced dimension space-time adaptive nulling. Nulling is achieved by adaptively combining non-pulse compressed data from a set of receiver elements. Three overlapped PRI sets are employed.

High Order Post-Doppler (HOPD)

The HOPD algorithm performs a reduced dimension space-time adaptive algorithm. Nulling is achieved by adaptively combining pulse compressed data from a set of receiver elements. No PRI staggering is employed.

The HOPD algorithm represents the most stressing of the three algorithms both in terms of

processor throughput and interconnect data rates. We have therefore chosen to address the HOPD requirements as part of the sizing effort for the real-time STAP processor. However, it is seen as an important requirement that the IOP design demonstrate the flexibility to perform the other algorithms as well. A flow diagram of the HOPD algorithm is shown in Figure 3.1.

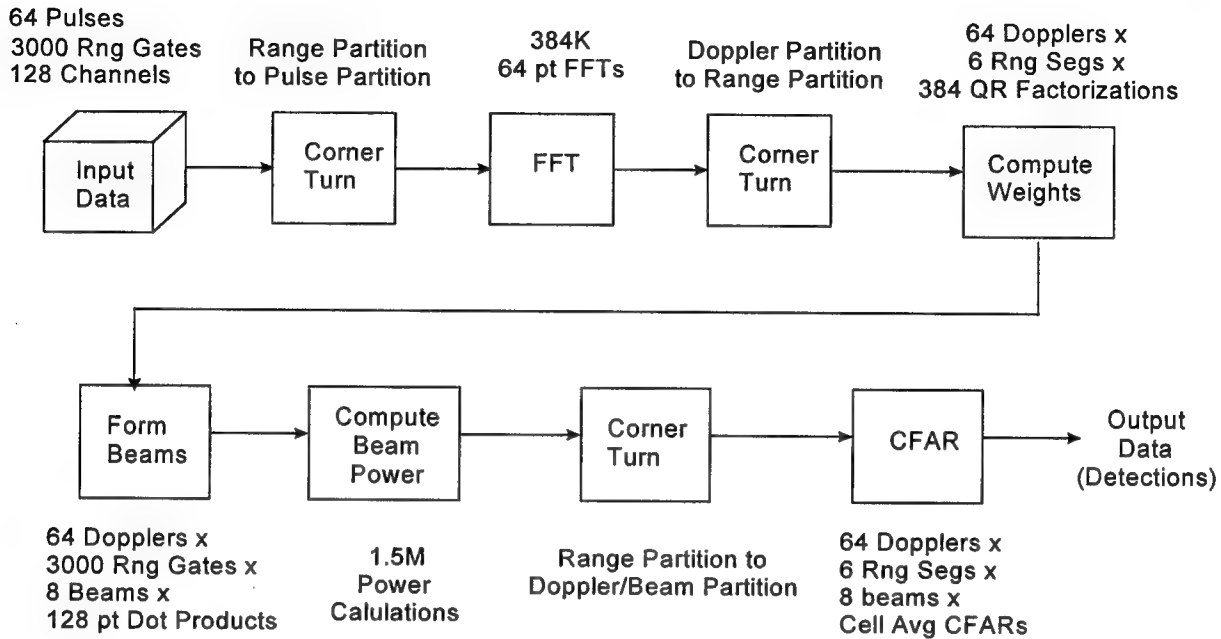


Figure 3.1 HOPD STAP Algorithm Flow Diagram

3.3 DEFINING IOP STAP SYSTEM PARAMETERS:

The Mountain Top program's STAP benchmarks specify a set a radar parameters applicable to each of the STAP algorithms. These are summarized in Table 3.2.

For the purposes of sizing a STAP processor as part of the IOP program, these parameter sets were found to be of limited value, since they do little to help bound the scope of the problem. *When taken as a whole they are found to span nearly 18 orders of magnitude.* At the low end of the Mountain Top spectrum, the STAP problem might be quite manageable using a commercially available processor, while at the high end, no technology known could meet the requirements.

To help focus our IOP efforts, a more limited set of STAP parameters was constructed. Where possible, these were based upon the proposed system parameters for the AWACS EOA UHF radar. This focused set of STAP parameters is also included in Table 3.2. Comparing our parameters with those supplied by Lincoln Labs shows that, in general, our targeted application falls somewhere within the mid-range of the parameters chosen for the Mountain Top program. The one notable exception occurs in the number of receiver channels, where it was felt that the Mountain Top parameter seemed to be a bit low. Ninety-six channels have been postulated for the AWACS EOA UHF antenna, therefore we have elected to address a 128 channel system to

allow for the potential of future growth.

TABLE 3.2 MOUNTAIN TOP/IOP STAP SYSTEM PARAMETERS

Parameter	Mountain Top Range	IOP Range	Units
Range Cell Rate	0.1 - 5.0	1.0	MHz
CPI Length	50 - 1000	200	msec
PRF	128 - 25000	333 / 667 / 1333	Hz
Receiver Channels	14 - 100	128	
Range Gates	84 - 20000	3000 / 1500 / 750	
Steering Vectors	2-32	8	
Range Segments	1-32	6 / 3 / 2	
Degrees of Freedom	42-600	128	
Training Samples	84-1200	500 / 500 / 375	
Doppler Bins	16-2048	64 / 128 / 256	

3.4 STAP PERFORMANCE GOALS:

In order to size a STAP processor based on the radar parameters defined above, some set of "success criteria" must be established. For a STAP processor, the rate at which adaptive weights can be generated generally serves to limit system performance. The more frequently weight sets are generated, the more accurately they serve to represent the clutter statistics currently being encountered by the radar.

For the HOPD algorithm, the Mountain Top program defines the success criteria shown in Table 3.3. As was the case with the radar parameters, the performance goals established for the Mountain Top program were seen to have limited value for use in defining the performance objectives of a flyable real-time STAP processor.

The Mountain Top end-to-end latency requirement of 3 seconds seems quite adequate for use in laboratory-based algorithm development, but was seen as impractical for use in real-time surveillance applications, particularly in the case of a beam-agile electronically steered array. Even if sufficient high speed buffer memory were available to store the raw radar data until the corresponding weights become available (this would require on the order of 12 Gigabytes of high speed RAM), the delay through the system would render it impractical for use in closed-loop tracking applications.

As a goal for the IOP program, the radar's coherent processing interval (CPI) was chosen as the minimum acceptable weight generation rate, while 0.5 seconds was chosen as the maximum acceptable weight generation latency for the STAP processor. The resulting system is thus amenable to use with an electronically steered array and to use in tracking applications.

As a goal, the size, weight, and power of the proposed IOP STAP accelerator were chosen to exhibit an order of magnitude improvement over those specified for the Mountain Top STAP processor. While this choice is admittedly somewhat arbitrary, the ability to claim an order of magnitude performance increase in parallel with an order of magnitude reduction in size, weight, and power was seen to serve as an important "selling point" when the potential technology insertion is eventually proposed to a potential customer.

TABLE 3.3 MOUNTAIN TOP / IOP STAP PERFORMANCE GOALS

Parameter	Mountain Top Goal	IOP Goal	Units
Weight Generation Rate	0.33	5	Hz
Weight Generation Latency	3.0	0.5	sec
Processor Size	20	2	ft ³
Processor Weight	600	60	lbs
Power Consumption	8.0	0.8	KVA

4. COTS-BASED STAP PROCESSING

Overview:

This section describes our assessment of the applicability of commercially available processors to fulfilling the IOP STAP requirements, and the resulting selection of proposed technology insertions into COTS-based architectures.

4.1 PERFORMANCE OF COTS-BASED STAP ARCHITECTURES:

Before undertaking the design of special-purpose STAP accelerator hardware, some attempt must be made to answer the question of whether or not such an accelerator is indeed necessary. Those proponents of "fully-programmable" COTS implementations would argue that the rapid rate of advance of commercially available processing technologies obviates the need for such a design.

In this regard, the results of the Mountain Top STAP study serve as useful metrics for assessing the ability of commercially available embedded processing architectures to meet our IOP STAP requirements. Inasmuch as these results are based upon the work of "independent" third parties (i.e. those not involved in the IOP program), they also serve as a more convincing, unbiased source of data for making such an assessment.

Figure 4.1-1 illustrates the achieved processing latencies for a variety of partitionings of the Mountain Top HOPD algorithm when hosted on a Touchstone Embedded Multi-computer and a Mercury Race Real Time Multi-processor. (These results are based upon work performed by Rome Labs, Honeywell, and Northrop-Grumman as part of the Mountain Top program). The solid lines represent data gathered on the Mountain Top program. These results show that the Embedded Touchstone processor was unable to meet the 3.0 second Mountain Top latency requirement, even with a total of 297 processors working together on the problem. The Mercury processor nearly met the 3.0 second latency requirement when 128 processors were employed.

An important characteristic observed in the measured results of Figure 4.1-1 is that the weight generation latency is decreasing nearly linearly with the number of processors assigned to the problem. This is due in large part to the natural parallelism in the STAP algorithm (with the important exception of the corner turn required as part of Doppler processing, which we will examine in more detail later in this report).

If we optimistically assume that this linearity can be extended *ad infinitum*, we may use the Mountain Top results to predict a lower bound on the number of processors needed to meet our IOP latency requirement of 500 msec. (The dotted lines of Figure 4.1 show the linear extrapolation of the achieved Mountain Top results.) With Mountain Top as a guide, we can thus predict that, at best, the IOP latency requirement is met when the number of processing elements grows to between 1000 and 3000.

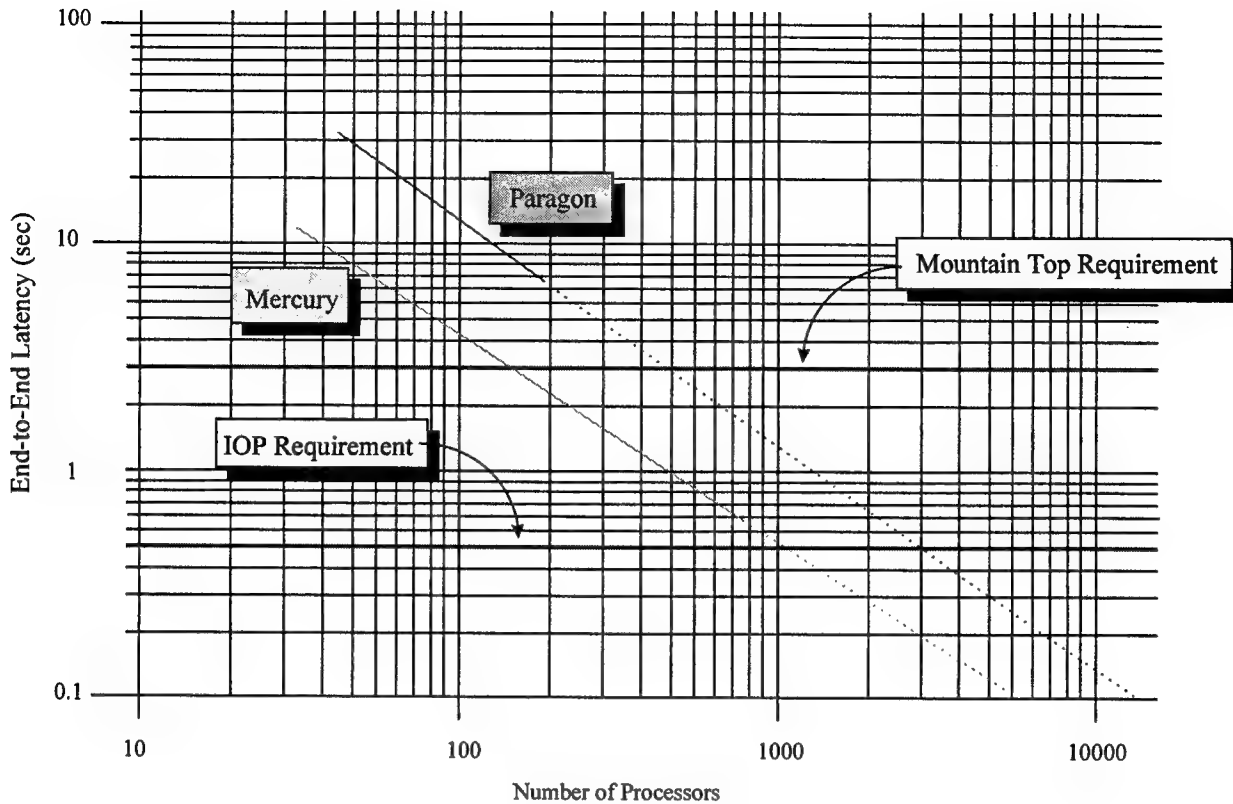


Figure 4.1-1 Achieved COTS STAP Processing Latency

These represent extremely large systems. (No such systems have ever been constructed.) The recurring cost of the system, its physical size and power requirements, along with the human programmer's ability to manage a task divided amongst so many processors bring into question whether such an approach is indeed even viable, especially given the constraints of an airborne application.

And it should be kept in mind that these numbers only address the system's end-to-end latency requirement. The additional need to sustain our "per CPI" weight generation rate (an issue which was not addressed as part of the Mountain Top program) will likely result in substantially greater increases in system complexity.

4.2 IDENTIFYING LIMITATIONS IN COTS-BASED STAP PROCESSORS:

Given the measured results from the Mountain Top program, it would seem that, at present, a "purely-COTS" based solution is ill-suited to application as a fly-able real-time STAP processor. However, given the rapid rate of advance of digital processing technologies, the nature of COTS-based STAP processing must be examined in further detail to determine whether the current limitations of COTS-based approaches to STAP will be overcome through evolutionary advances to the state of the art.

In an effort to assess the apparent limitations of COTS based approaches to STAP, the measured Mountain Top results again provide useful insight.

To characterize the ability of a multiprocessing system to carry out a distributed computation, it is useful to define the "speedup" of the computation, which may be defined as the ratio of the time required to perform the computation on a single processing element to that required to perform the same computation on a collection of N processing elements, i.e.

$$S_N = T_{E1} / T_{EN}$$

In theory, for a perfectly efficient distributed process partitioning, achieved speedup should approach N. It should also be noted that fractional speedups (i.e. speedups less than unity) are possible. Such a result implies an operation which requires more time to execute when distributed over a collection of N processors than on a single processor.

The achieved speedups for the three Mountain Top STAP algorithms, as measured on a Mercury Race processor, are illustrated in Figure 4.2-1. Here the algorithms have been subdivided into their five principal constituent sub-functions, namely: corner turn, FFT, assemble QR data, compute weights, and apply weights. To avoid cluttering the graph, individual sub-functions have not been shown, however, the sub-functions have been segregated into "compute" tasks (FFT, compute weights, apply weights, and "communications" tasks (corner turn, QR data distribution).




Figure 4.2-1 Achieved Speed Up of Various STAP Partitionings

These results are enlightening. They show that in all cases, the "compute" functions achieved nearly the theoretical best possible speedup when distributed across multiple processors. This speedup was maintained down to a quite fine-grained process partitioning, and is due in large part to the natural parallelism inherent in the STAP problem.

The "data transfer" functions, by contrast exhibited a speedup of less than unity, i.e. these functions consumed more and more time as the number of processing elements was increased.

This is an important result. *It clearly points away from electronic processing, and toward data communications, as being the key limiting factor in the performance of a real-time STAP architecture.*

4.3 QUANTIFYING THE IMPACT OF DATA COMMUNICATIONS IN COTS-BASED STAP:

By analyzing measured results from the Mountain Top program, we have identified data communications as being the key factor limiting the performance of a COTS-based STAP processor. Next, we will attempt to quantify this performance impact.

To that end, Figure 4.3-1 illustrates the measured cumulative contributions of communications and processing, as a percentage of total execution time for a variety of partitionings of the three Mountain Top STAP algorithms, measured on a Mercury Race computer.




Figure 4.3-1 Contributions to STAP Execution Time of Constituent Sub-functions

These results show that once the STAP computations grow to a point where 100 or more processors are required, data communications does indeed account for a significant percentage of the overall execution time of the STAP algorithm. In fact, in two of the three cases, once the algorithm partitioning becomes sufficiently fine-grained, data communications accounts for the majority of the end-to-end latency through the STAP processor.

It must be kept in mind that these results are based on a much simpler problem (3 sec allowable latency) partitioned onto at most a few hundred processors. The targeted IOP requirement of a 500 msec latency is estimated to require 1000 or more processors. Thus, given the trends predicted here, the IOP problem will likely become completely I/O bound before the necessary number of processing elements can be brought to bear on the problem.

4.4 LOCATING THE COTS DATA COMMUNICATIONS BOTTLENECK:

To better illustrate the difficulties encountered when trying to map the STAP problem into a COTS environment, it is useful to refer back to the algorithm flow diagram of the HOPD STAP algorithm of Figure 3.1. A number of "corner turn" operations are shown in the data flow diagram. These transformations essentially perform a matrix transpose operation, and are an essential part of all Doppler radar processing (whether adaptive or not). They serve to reorder the data from a range-order (the order in which the radar signal naturally returns to the receiver) to the pulse-order required to perform the Doppler FFT. The post-FFT'd data (now in Doppler-order) is subsequently converted back to range-order for the range segmented STAP processing. Later, yet another corner turn reorders the data back into Doppler-order prior to CFAR.

This corner-turn process cannot, to any great extent, be pipelined. The transfer cannot commence until all range-ordered data has been received, and no further processing can begin until the last pulse-ordered data has been exchanged amongst the processors.

Since these corner turn operations essentially require "zero" processor operations, their impact can be completely lost in the traditional "ops count" used in processor sizing estimates, and yet we have found them to be the principal factor limiting our ability to perform real-time STAP in a COTS environment.

The difficulty in performing such as seemingly simple operation in a parallel processing environment stems from the fact that the data to be transposed are physically distributed amongst a number of the system's processing elements (distribution amongst 10's or 100's of processors would not be uncommon). The transpose operation thus requires the simultaneous transfer of many small amounts of data between large numbers of processors. This "scatter-gather" data flow is shown in Figure 4.4 which illustrates pictorially, a small distributed corner turn.




Figure 4.4 A Distributed Corner Turn

A variety of factors contribute to poor corner turn performance in a distributed system. These include:

- Contention: large numbers of processors simultaneously accessing the system's shared interconnect resource generate significant bus contention, and "wait states" as some transfers are blocked, and denied access to the bus.
- Overhead: each processor must carry out a large number of transfers each of which is uncharacteristically small in size. The overhead involved in gaining access to the system's interconnect resource thus constitutes a substantial percentage of the time of each transfer.
- Page faults: The data to be sent by each processor is stored non-sequentially in its local DRAM memory. This non-sequential access results in DRAM page "thrashing", whereby the access-time penalty paid for crossing a page boundary of the DRAM devices (which usually occurs infrequently during more typical sequential accesses) limits the achievable data memory access rate.

4.5 SOLUTIONS TO THE DATA COMMUNICATIONS BOTTLENECK:

Several enhancements to existing COTS architectures may be used to reduce the effects of this data communications bottleneck, making possible cost effective real-time STAP processing. We propose to employ each of these approaches in the design of the IOP STAP processor.

4.5.1 Resource Concentration (Hardware Acceleration):

We have shown the scatter-gather behavior of distributed STAP processing to be problematic. One obvious solution then is to eliminate the need for scatter-gather data flow.

By enhancing our existing COTS system with a specific set of optimized resources (i.e. hardware accelerators), processing throughput may be concentrated into a smaller physical space (e.g. a single circuit board or set of boards). In this way, we eliminate the need for the scatter-gather of data between physically (or logically) distant processing nodes via the system's shared resource fabric. Localizing data flow provides fuller, faster, and less expensive connectivity.

However, with special purpose designs come concerns; most notably *cost* and *lack of flexibility*. Non-recurring development costs (especially when ASIC designs are involved) can be high. Since the specialized nature of the design may limit its widespread applicability, low volume production runs mean higher recurring cost per board than a more general-purpose design. Since the designs have been optimized to perform a specific function, they may be unable to adapt to future algorithmic change.

Nonetheless, the high payoff for developing a hardware accelerator is often warranted. In our IOP design for example, we have been able to replace a thousand or more general purpose processors with a single circuit card. The development cost of the accelerator may fall well below the software development cost involved in hosting a problem onto a 1000+ processor system. And since so few boards are required, the recurring cost of the accelerator will be much less than that of the many processor cards required by the general purpose approach.

In our IOP design, candidates for hardware acceleration have been carefully chosen so as to maximize their general applicability. By identifying and accelerating a set of STAP "building block" functions, rather than a specific STAP algorithm itself, our chosen accelerators are useful in a wide variety of STAP applications (e.g. they are capable of performing all three Mountain Top STAP algorithms). They are parametrically programmable (number of receiver channels, degrees of freedom, range cells, etc.), and modularly constructed, so that they may be cascaded in either dimension to solve larger STAP problems.

Our accelerator tradeoffs, selections and designs are fully documented in Section 7 of this report.

4.5.2 Increasing Architectural "Robustness" (Topology):

From a performance perspective, the best architectural topology for any distributed processing system is a fully connected network, where all nodes can communicate simultaneously. However, such a topology is usually impractical in larger systems. Compromises such as *hypercubes*, *meshes*, and *fat trees* result.

The amount of connectivity (as opposed to raw link bandwidth) is often the key limiting factor in the ability to distribute processes algorithmically. Thus, increased connectivity is highly desired.

The RaceWay interconnect fabric employed in the Mercury family of real-time multiprocessors is uniquely suited to benefit from increased interconnectivity. Based on a partially blocking crossbar scheme, known as a "*fat tree*", the RaceWay places no architectural limit upon the number of direct interconnect paths which may enter or exit a processor (in contrast to meshes, hypercubes, etc. which impose strictly defined architectural limits). Thus, the RaceWay is able to take direct advantage of any technology which allows greater connectivity.

Connector density is key here. Today's Mercury RaceWay uses 40 parallel wires to pass data via each channel. This number of wires, given the inherent electrical constraints on pin density and power consumption, limits RACE systems to about 4 data channels off each 6U board. This 4 channel limit is why RACE systems often use active backplanes.

As part of our IOP design, we propose to develop a new hybrid electro-optical RaceWay crossbar "chip" which integrates both electrical and optical ports. The parallel electrical ports will serve processors co-located on the same board. The serial optical ports will pass data between boards and chassis. The hybrid crossbar chip must use a very small optical connector, permitting users to squeeze many ports onto a card edge. Such an electro-optical crossbar would enable construction of topologies that are not possible today, allowing the user to adapt the topology to the application; typically much easier than trying to recast an application to fit a fixed topology.

Our work on modifying the existing Mercury architecture is documented in Section 5 of this report. Details of our work on the new electro-optic RaceWay crossbar chip are contained in Section 6.

4.5.3 Increased Link Bandwidth:

Historically, the bandwidth between any two points in a system area network (SAN) has been matched to the bandwidth between a processor and its DRAM. One reason for this "rate matching" is that it was very expensive (in terms of design complexity and silicon) to create rate change buffers between the memory and the communications link. The Mercury RaceWay operates at bandwidths of 160 Mbytes/sec, a common maximum for DRAM technologies in 1993.

As memory access speeds increase, one must consider the cost of increasing the SAN link bandwidth to keep place. As a result, SAN data rates (including the RaceWay) have begun to lag behind memory access rates (estimated to reach 2 Gigabytes/sec by the year 2000), and rate matching buffers have become commonplace.

The RACE processor family will eventually move beyond 160 Mbyte/sec. Two projects with such a goal exist. RACE 2.0 will leverage serial, electrical innovations like LVDS (Low Voltage Differential Signaling). RACE 3.0 will use optics.

Detailed schedules for RACE 2.0 and 3.0 do not exist. Mercury Computers is waiting for the semiconductor community to deliver appropriate link technologies in production volumes. If the optics community delivers first, plans for an intermediate RACE 2.0 will be dropped. Any link technology RACE moves to must be cost effective. This generally means that the technology must have a volume driver outside the embedded computing niche.

4.6 THE NEED FOR OPTICS

We have identified the key limitations which stand in the way of a COTS-based STAP processor, and have identified several means through which we may extend the capabilities of COTS processors to make STAP processing viable. Here, we illustrate the need for optics to bring about these innovations.

4.6.1 STAP Accelerator:

Application specific accelerators allow the concentration of computing resources, thereby localizing much of the required communications to within a small set of boards, removing it from the shared interconnect fabric which has been identified as the key STAP bottleneck.

However, this use of accelerators has a secondary effect. Concentrating compute resources also means concentrating data flow. Data flowing between accelerator cards exceeds the capacity of the COTS system's existing interconnect media (e.g. the Mercury RaceWay).

Our reliance upon accelerators demands the use of a physically compact, high speed, narrow word-width interconnection media. The high speed and small physical footprint of optical interconnects are an integral part of our accelerator design approach.

4.6.2 System Topology:

More generous connectivity is essential to the development of fuller, more robust system topologies, with the goal of eliminating the view of the system's interconnect structure from the programmer. Currently, physical constraints of connector pinout and electrical constraints on power consumption limit the amount of board to backplane connectivity which may be provided.

High density board-edge connectors are the key innovation needed to allow increased interboard connectivity. Our IOP design relies heavily upon high density optical connectors, and upon the development of an electro-optic crossbar switch to achieve the increased connectivity needed to make STAP processing viable.

4.6.3 Passive Backplanes:

The constraints of routing density limit the amount of signaling on the system backplane. To meet the demand for global system inter-connectivity, partial distributed crossbars, with active backplanes (i.e. active circuitry mounted as part of the backplane) are the result.

Active backplanes are a prime maintainability concern in military systems. (A failed circuit card can be replaced in flight while the failure of backplane circuitry requires cancellation of a mission.)

The speed of optical interconnects allow communications paths to be highly serialized, and carried via just one or a few interconnects (versus the 40 interconnects needed per communications path today). This speed, along with the high density of optical waveguide backplanes makes possible the use of a centralized switching network, comprised of a single removable (and hence replaceable) circuit card, in place of today's distributed active backplane circuitry.

Not only does this approach eliminate the need for an active backplane, but it allows for significantly fuller, non-blocking connectivity through the centralized switching fabric.

Our IOP design relies heavily upon centralized switching and optical backplanes to deliver the increased connectivity needed to achieve real-time STAP performance.

4.6.4 Sensor-to-Processor Interconnects:

Data from hundreds of receivers must be collected and funneled down into a small set of STAP accelerator boards. This implies the need for a high speed, narrow word-width communications means. Due to platform constraints, these interconnects must be treated as physically distributed, covering tens of meters, and being subject to a harsh electrical environment (EMI).

Electronic interconnects are ill-suited here. High speed, light weight, interference resistant optical interconnects are a key component of our multi-channel STAP system.

4.7 SUMMARY

Through the analysis of existing measured data derived from the Mountain Top program, the inherent limitations of COTS-based architectures to the successful implementation of STAP processing in an airborne environment have been identified and quantified. Data communications has been shown to be the principal limiting factor in the ability to achieve real-time STAP performance via a COTS-based multiprocessor.

A variety of technology insertions based upon available, laboratory proven technologies have been proposed as means to overcome these limitations. The work on each of these insertions is documented in the subsequent sections of this report.

5. IOP ARCHITECTURE

We have selected a Mercury RACE-based COTS architecture as the foundation for the IOP. We believe the RACE architecture to be uniquely suited to exploiting the benefits of innovations in interconnect technologies such as optics.

We have incorporated several key interconnect technologies being developed by the OMNET and POINT programs into the RACE architecture. In addition, we have augmented the performance of a "COTS-only" solution through the integration of a set of RACE-compatible STAP accelerator modules.

The resulting IOP design is thus amenable to fabrication in extremely small size, weight, and power, and to a low cost of manufacture. In this section, we examine our IOP architecture and the factors which drove its selection.

5.1 SELECTION OF THE MERCURY RACE AS A BASELINE STAP ARCHITECTURE

The development of complex systems like AWACS tends to be evolutionary rather than revolutionary. With the demand for thirty to forty year (or longer) platform lifecycles, the design of an IOP in the 1990's must be done in the context of technologies of the year 2000, and well beyond, if the system is expected to remain viable and useful (by means of evolutionary upgrades) throughout its lifetime.

Thus, the ability of a system architecture to exploit new (and perhaps radically different) technologies, which emerge during the course of its lifetime, becomes a fundamental issue in the viability of its design.

We see this capability as the key strength of the Mercury architecture. Based on a partially blocking crossbar scheme, known as a "*fat tree*", the RACEway interconnect fabric employed in the Mercury processor family places no architectural limit upon the number of interconnect paths which may enter or exit a processor. This is in contrast to more rigorously "*well defined*" structures such as meshes, hypercubes, etc. which impose strict architectural limits.

As an architecture then, the Mercury RACE is uniquely suited to benefit from any emerging technology which allows increased interconnectivity. This is a clear advantage for a STAP processor, since we have identified the need for increased connectivity as being the fundamental issue in our ability to achieve real-time STAP performance.

5.2 ARCHITECTURAL DECISIONS DRIVEN BY THE INTENDED APPLICATION OF THE IOP

The inherent characteristics of STAP serve as key discriminators driving the selection of the IOP architecture. Figure 5.2-1 illustrates a generalized characterization of the STAP problem.

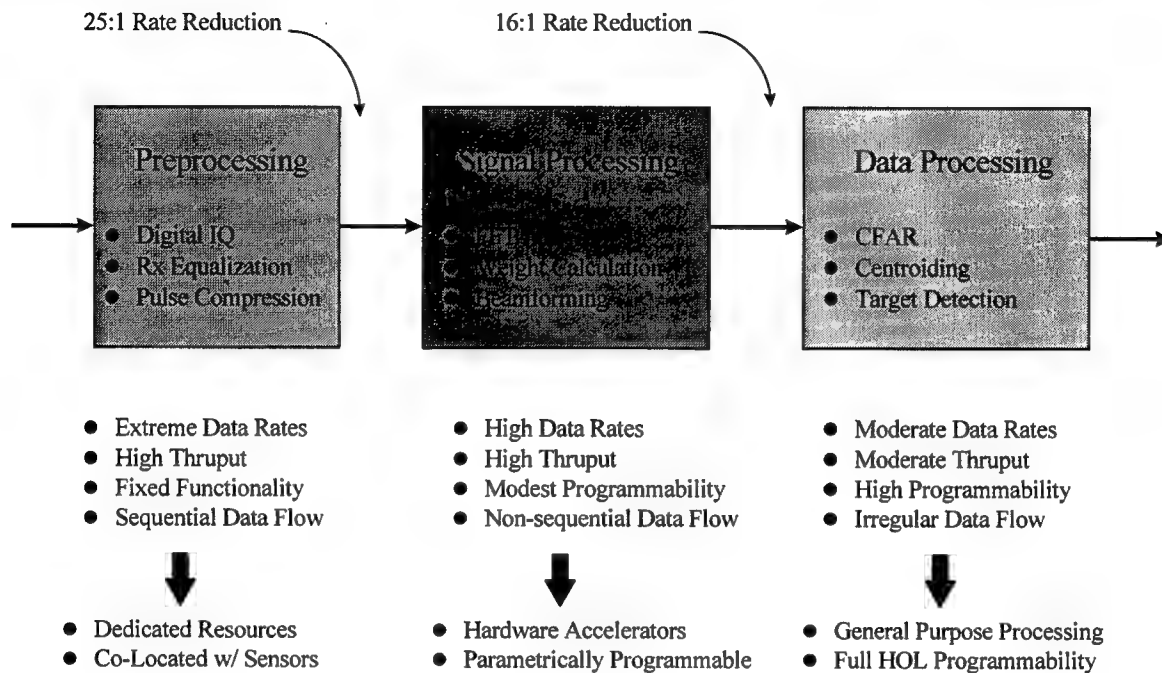


Figure 5.2-1. STAP Processing Characterized

Front-end *preprocessing* of raw sensor data is characterized by demanding throughputs, fixed functionality, and no data dependency. As a result, little or no benefit is derived from employing programmable resources here. The demanding data rates flowing between the sensor and preprocessor, coupled with the potential to exploit a 25:1 data rate reduction following preprocessing argue persuasively that these functions should be physically divorced from the central STAP processor, being instead geographically co-located with the sensor.

Signal processing, in essence the “extraction” of discernible information from the radar signal, is characterized by high throughputs, stressing regular but non-sequential data flow, and the need for modest (i.e. parametric) flexibility. It is here that we have identified a need to reduce the burden on the shared interconnect structures of “COTS-only” approaches to STAP. Thus, our architecture will employ hardware accelerators and the increased connectivity afforded through the use of optical interconnects as potential means to solve this data communications bottleneck.

Data processing, which we classify here as the “analysis” of the information extracted during signal processing is characterized by relatively modest throughput and data rates, but significant data dependency and irregular data flow. This processing is a natural candidate for the application of fully programmable, general purpose processors communicating via a shared global interconnect structure.

5.3 ARCHITECTURAL DECISIONS DRIVEN BY THE OPERATING ENVIRONMENT OF THE IOP

Considering processing throughputs and data rates alone is insufficient for selecting an architecture appropriate for the IOP. We must also examine the environment, both “physical” and “economic” in which the IOP must operate.

5.3.1 Airborne Environment:

The limitations of an airborne environment place significant constraints upon the physical mechanization of our IOP design. For example, it could be argued that incorporating a more flexible optical interconnect structure alone enables the STAP problem to be solved (via a 1000+ node system). Such a solution, however, ignores the size, weight, power, and cost constraints of our targeted application. Our approach must specifically address both the need for COTS flexibility as well as the physical and cost constraints associated with an airborne system in volume production. Our approach must not only makes STAP “possible”, it must make it “practical”.

5.3.2 Software:

Escalating development costs are the key issue driving the move toward COTS-based military systems. COTS suppliers, such as Mercury, are able to amortize development expenses over multiple programs, holding down costs for each customer. (More than half of Mercury’s research and development investments underwrite the software infrastructure users need to program in a massively parallel environment.)

Compatibility with this software infrastructure is a fundamental goal for our IOP design. A development program such as AWACS-EOA cannot afford to create a similar infrastructure from scratch. Our architecture is designed to exploit the flexibility afforded by this existing software environment, while simultaneously bolstering the efficiency of conventional COTS processors by incorporating application specific accelerators and a more flexible optical interconnection scheme.

Our optical interconnect will have the same routing scheme used in today’s RACEway. Thus system software only needs to change in reaction to the improved topology (topology changes impact software routing strategies). This is a software area that Mercury, not its customers, must adapt. For this reason, Mercury carefully designed its COTS software to make topology-oriented enhancements easy. The system software changes to support the hardware defined in this study should require less than \$50,000 in NRE.

Some application software may also require enhancement. We refer here to application software that defines explicit paths through RACEway. Such software must change whenever the underlying hardware changes.

5.3.3 “Core technology” vs. “Off-The-Shelf”

In today’s cost-conscious environment, military end-users are no longer willing to foot the bill for the development and insertion of emerging core technologies into systems. The cost and political risk to full-scale development programs is simply too great.

And yet, we must look to alternative solutions if we are to provide STAP capability to an airborne system. Thus, as a key element of our IOP design, we plan to exploit core technologies to be developed under the OMNET, POINT programs. We intend to use IOP as a means to mature the outputs of these core technology demonstration efforts to the stage where they will be ready for system insertion.

5.3.4 Cooling

Whenever designers squeeze increasing functionality into a constrained space, power usage per square inch tends to go up.

There are two challenges associated with high power VMEbus boards. The first challenge involves getting power into the board (providing enough power pins at proper voltage levels). The second and greater challenge involves radiating power off a VMEbus board.

Our team’s principal answer to both challenges is through heavy use of DSP technology. DSP chips provide significantly more performance per watt than conventional microprocessors. Our team leverages this fact to avoid making exotic packaging innovations.

To achieve the above, we will seek to keep power requirements down to about 30 watts per 6U VMEbus card. Succeeding here requires next generation DSP components (so we don’t need quite so many DSP chips on a single card). Our team will also spread the weight maker ASICs over more VMEbus cards than we initially expected. This decrease in weight maker density keeps us within air-cooling power density requirements.

5.4 KEY INNOVATIONS REQUIRED FOR IOP

We view the contribution of optics toward solving the limitations on architectural topology imposed by electronics as more important than increasing the raw bandwidth of individual links. *To be used to full advantage, optics must be viewed as more than just a high speed replacement for copper interconnects.*


An essential building block needed for IOP is a hybrid opto-electronic router (crossbar switch), which exploits the interconnect density inherent to optics to increase the connectivity onto and off of a circuit card. Currently, we envision four electrical ports and 5 optical ports (as a

minimum). The electrical ports will use parallel signaling and serve on-board processors. The optical ports will serve processors between boards and chassis'. (We do not care whether the optical ports use parallel or electrical signaling, only that a large number of ports fit onto a card edge, that the optical ports consume very little power, and that the system be affordable).

DARPA's OMNET program will develop an opto-electronic switch. However, it will likely not achieve this desired port density. In addition, it will employ fiber ribbon interconnects. Our selected approach will build upon what is learned via the OMNET research prototype to construct a higher density electro-optical switch using a polymer optical backplane to deliver the interconnect density desired for IOP (see Figure 5.4-1).

OMNET

12 channel fiber ribbon
MT connectors
34 signals per card edge inch



IOP

Polymer optical backplane
Passive alignment connectors
250 signals per card edge inch




Figure 5.4-1 IOP Advances to OMNET Optical Backplane Technology

5.5 CONCLUSIONS:

Our selected IOP architecture is illustrated in Figure 5.5-1. We have chosen to combine a set of application specific STAP accelerators and general-purpose high performance computing (HPC) nodes, as part of an overall programmable COTS processor, the system being integrated via a homogeneous optical interconnect structure based upon the ANSI standard RACEway.

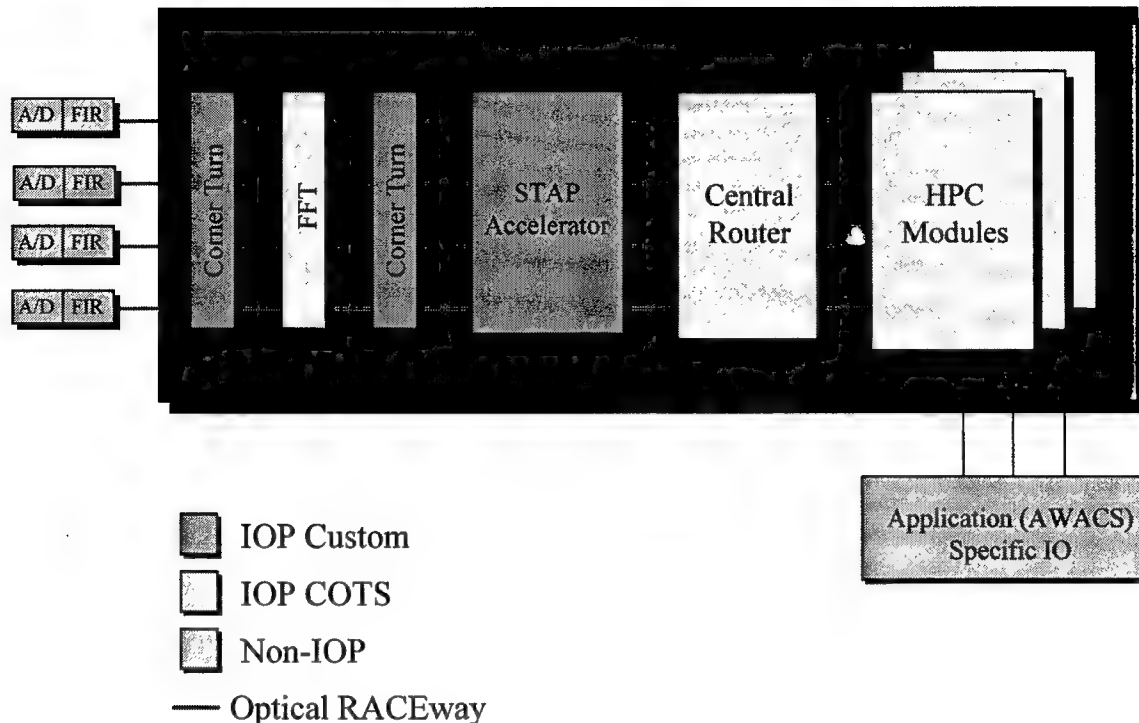


Figure 5.5-1 IOP Mercury-Based STAP Architecture

We see the following distinct advantages to this approach:

The design is highly efficient. By integrating a set of application specific accelerator modules into a COTS-based computing environment, we are able to provide the high processing throughputs demanded by STAP applications, without sacrificing the flexibility afforded by a COTS software environment.

Our design allows the STAP processor to reside within a single 6U VME chassis. We have thus opened up the potential for incorporating STAP into a wide range of applications where size, weight, and power constraints have heretofore rendered it unfeasible, including airborne surveillance, UAV's, and airships.

Our design leverages and advances key innovations from DARPA's POINT (polymer waveguides) and OMNET (opto-electronics and packaging innovations) programs to provide significantly increases in interprocessor communications. We thus eliminate the cost and risk associated with the insertion of unproven core technologies into systems.

Our Mercury RACE-based system is compatible with a large volume of existing software. This provides an upgrade path for significant airborne surveillance applications.

Our design replaces over 1000 HPC modules with a set of two accelerator module designs. Thus, our approach is highly reliable, is easy to program, and is affordable enough for use in volume production.

Our approach eliminates the need for an active backplane, and is more maintainable than Mercury's existing COTS product line. The IOP design leverages optics to eliminate the need for active backplane circuitry while simultaneously delivering critical bandwidth-density advances. All active components in our proposed configuration are provided as easily replaced modules. Hot swapping is possible in theory, however, Mercury COTS software does not yet support live reconfiguration but could be made to do so in the future.

6. OPTICAL INTERCONNECT TECHNOLOGY

For many years, Honeywell has been an industry leader, and made tremendous advances, in developing photonics components, serial and parallel transmitter (Tx) and receiver (Rx) module packaging technologies, and complete links at both the inter- and intra-cabinet levels. We believe that these recent developments have made it possible, and practical, to insert completely transparent digital-to-digital optical links in real systems such as the IOP.

6.1 OPTICAL LINK COMPONENTS AND BUILDING BLOCKS

6.1.1 Vertical cavity surface emitting lasers (VCSELs)

VCSELs have emerged as a powerful new device technology. First developed in laboratories, they have quickly matured to become a commercial product in just a few years. Honeywell has been in the forefront of both research and product development of short wavelength (850 nm) VCSELs for data communications.. Compared to conventional edge emitting lasers, VCSELs have the advantages of lower threshold current (a few to sub-mA),




Figure 6.1-1: 32-channel VCSEL array

very high speed (3dB bandwidth of over 14 GHz has been measured), less temperature sensitivity of threshold current, wavelength shift, and output power (an operating range of -55C to 125C has been demonstrated), easy wafer level batch testing, symmetric and less divergent output beams which facilitate fiber coupling and packaging. In addition, VCSELs are easy to fabricate into 1D (see Figure 6.1-1) or 2D arrays which are critical for parallel optical interconnects. Honeywell is the first company that has offered VCSELs as a qualified (with complete reliability tests) commercial product for optical data communication. The VCSELs have become a key component for state-of-the-art high speed digital optical links. Not only is the reliability excellent for commercial applications, but initial data suggests that they will be well suited for military applications.

6.1.2 Integrated optical receiver

Optical receivers (Rx's) with integrated photodetector and electronic amplifiers are another key element for high speed optical links. Compared to hybrid detector and amplifier, integrated receivers have the advantage of less parasitic capacitance (and therefore better noise characteristics at high speed), allowing simple and low cost packaging, and being easily scaleable into arrays for parallel links.

For 850 nm wavelength data links, GaAs metal-semiconductor-metal (MSM) photo-detectors are compatible with widely available high speed GaAs MESFET technology. Over the past years Honeywell as well as other companies have been developing various type of high speed integrated receivers and arrays (see Figures 6.1-2, and 6.1-3) ranging from analog to digital and from several hundreds Mbps to several Gbps, all fabricated through commercial GaAs foundries. As a sign of wide acceptance of the technology, such an integrated GaAs optical receiver at 1 Gbps has recently become a commercial off the shelf (COTS) product offered by manufactures such as VITESSE Semiconductors. Figure 6.1-4 shows a eye diagram measured for a link consisting of Honeywell VCSEL and a VITESSE made COTS Rx.

6.1.3 Polymer optical waveguides:

For intra-cabinet optical interconnect media, Honeywell has been developing polymer optical waveguide technology that is equivalent in function, and compatible in material, fabrication, and assembly processes with existing electronic printed wiring boards (PWB) manufacturing infrastructures. Honeywell has used polyetherimide (Ultem from GE) and benzocyclobutene (BCB, from Dow Chemical) as waveguide core and cladding materials, respectively. Both polymer materials have proven stability and reliability over wide temperature ranges, and are widely used in microelectronics and other industries. The waveguides are normally multimode (see Figure 6.1-5(a)) for easy fabrication, handling and alignment. They are typically made on flexible substrates such as Kapton sheets (see Figure 6.1-5(b)), which can then be laminated into multi-layer boards through a standard board lamination process. There are also several other polymer technologies such as Dupont's Polyguide, and Allied Signal's Acrylate monomer based waveguides which are also very promising for




Figure 6.1-2: GaAs foundry processed integrated Rx layout, chip, and package




Figure 6.1-3: 12-ch. integrated Rx design for 2.5 Gbps per each channel.




Figure 6.1-4. Eye diagram at 1.6 Gbps Honeywell VCSEL → GaAs COTS Rx

optical backplane interconnects, although data on performance in aerospace environments is not yet available.



Figure 6.1-5. (a) Cross-section microphotograph, and (b) photo of ULTEM/BCB polymer waveguide fabricated on Kapton substrate.

6.1.4 Waveguide connectors

A optical link for intra-cabinet interconnect is not complete without proper waveguide connectors at MCM-to-board and board-to-backplane boundaries. Honeywell believes that viable solutions for connectors of this kind must offer two key features: (a) high density, and (b) manageable alignment tolerance. Honeywell has developed a MCM-to-board connector based on *passive* alignment of flexible waveguide arrays to board waveguides using a keeper piece (see Figure 6.1-6).




Figure 6.1-6: Waveguide MCM to board connector with passive alignment feature

Honeywell has also demonstrated a board-to-backplane connector concept based on an expanded-beam approach that connects arrays of waveguides (over 30) using a single pair of 3 mm graded index lenses (see Figure 6.1-7). This approach provides sufficient alignment tolerances to allow such optical connectors to be implemented inside existing electrical connector housings. Honeywell is now working with connector manufacturers such as AMP to implement such a concept using binary optical lenses, instead of costly micro-optical elements.




Figure 6.1-7: Board-to-backplane connect based on expanded beam approach.

6.1.5 Tx/Rx Packaging, and optical links

Honeywell believes that optical interconnect technology has to provide a transparent digital-to-digital solution to the system user. This solution can be realized through packaging optoelectronic components (such as VCSELs, photodetectors or integrated receivers) with other necessary electronics ICs (such as laser drivers, receiver amplifiers, Mux/Demux and Code/Decode ICs) into proper optical transmitter (Tx) and receiver (Rx) modules. These packaging technologies and resulting modules provide proper interfaces with both electronic logic and optical media such as optical fibers or waveguides, allowing system integrators to use the technology to solve real world problems.

The packaging technologies for optical data links available today include already standardized fiber based serial data communication Tx/Rx modules, such as Honeywell's gigabit modules. More advanced packaging includes parallel fiber ribbon based modules such as the 32-channel modules (see Figure 6.1-8) developed under the Optoelectronics Technology Consortium (OETC) by Honeywell, IBM, AT&T, and Martin-Marietta. Honeywell has also developed a gigabit link module (Figure 6.1-9a) and a high speed fiber optical data bus (FODB) 3-channel Tx/Rx multi-chip module (MCM) for satellite applications (Figure 6.1-9b). Honeywell has also developed advanced packaging technologies for both intra-cabinet optical waveguide-based parallel optical Tx/Rx modules using conventional electronic MCM-ceramic and some more advanced MCM like GE's HDI packaging (Figure 6.1-9c,d). The packaging efforts includes unique device-to-waveguide interface coupling technologies that allow passive alignment of waveguides to VCSELs and photodetectors.





Figure 6.1-8: OETC 32-channel parallel fiber link Tx modules



**Figure 6.1-9: (a) Gigabit Tx/Rx Module; (b) Three-channel FODB Tx/Rx MCM
(c) HDI packaged Tx/Rx array with connectorized polymer waveguides;
(d) MCM-Ceramic packaged multi-channel Tx/Rx with polymer waveguide interfaces.**

6.2 OPTICAL DATA LINK

6.2.1 Serial vs. Parallel:

The tradeoff between serial and parallel optical links is determined mainly by the system bandwidth requirement and implementation cost. In Figure 6.2-1, we show a typical generic optical link diagram. At the input of the Tx is a parallel N -bit wide electrical data bus operating at a low speed of B bits per second. In a typical serial link one uses $m:1$ Mux at transmitter (Tx) to “serialize” the parallel data and reduce the signal to N/m channels (often $N/m=1$), and increase the data rate of each optical channel to higher $m*B$ bits per second. The maximum serialized bandwidth, $m*B$, expected in the near future will be 1~3 Gbps without resorting to very expensive telecommunication technology. As the system data throughput increases, the number of high speed optical data channels will increase to the point the extra cost and space required will justify the use of parallel optical links where $N/m \gg 1$ channels are packaged into a single compact module with one or two ribbon (fiber or waveguide) connectors. In some applications, particularly in intra-cabinet interconnects, parallel interconnects are desirable or required as one can not afford the extra space, power, and signal latency introduced by the Mux/Demux chips.




Figure 6.2-1: Schematic block diagram of an optical link


6.2.2 Fiber vs. Waveguide:

In replacing co-axial cables between systems (inter-cabinets and sub-networks), optical fiber works extremely well for point-to-point optical data communication links over distances from a few meters to a few hundred of meters. We believe, however, that optical waveguides will be much more practical in replacing signal lines on PWBs between boards or backplane, and modules (intra-cabinet) where space is at premium and a large number interconnects are required at very high density.

Advantages of fiber based optical interconnects (both serial and parallel) include highly scaleable link bandwidth and distance, light weight, low power/distance, and low electromagnetic interference (EMI). Figure 6.2-2(a) shows a parallel 12-line fiber ribbon terminated with an MT connector. This fiber ribbon has a standard center-to-center pitch of 250 μm .



(a) 12-channel optical fiber ribbon with MT connector



(b) 144-channel optical waveguide ribbon with connector

Figure 6.2-2: Optical fiber and waveguides for parallel optical interconnects

For polymer waveguide-based intra-cabinet applications the most obvious advantage of using optical interconnects, in addition to those mentioned above, is that it offers higher interconnect density through a boundaries of packages, and boards-to-backplane. Figure 6.2-2(b) shows a 144-channel polymer waveguide ribbon of terminated with a MT connector. This optical waveguide has a center-to-center pitch of 100 μm . In the next section, we will describe some special design considerations of optical links in the context of the IOP program.

6.3 OPTICAL INTERCONNECTS FOR IOP MERCURY SYSTEMS

The objective of the IOP program is to use optical technology to improve the processor performance of the next generation AWACS. As an optical technology provider, Honeywell has been working with other team members including Northrop Grumman, Mercury Computing Systems, and Syracuse Research Corporation to understand the specific system data throughput requirements and performance constraints, and to offer solutions to the interconnect problems identified. Using optical backplane technology, the IOP STAP system can be implemented in a

single 6U VME chassis without requiring the use of active backplanes, considered undesirable for reasons of maintenance.

6.3.1 The Mercury RACE Architecture

Mercury Computing Systems' high performance parallel computers are built on a switching network based on RACE technology. The RACE technology has become the emerging standard for system area network adopted by more than a dozen system vendors. The building block of Mercury's switching network topology is the RACE Crossbar chip (CYPRESS part # CY7C965). The RACE crossbar chip has six bi-directional ports, each port having a bandwidth of 160 MB/s or 1.28 Gbps (32-bit wide data lines clocked at 40 MHz).

Current Mercury systems link each processor through a "fat tree" type network topology which consists of multiple levels of the crossbar switches. Figure 6.3-1 shows a small 30-node system, which consists of eight VME boards, connected by an eight-slot backplane. The limitation of board edge connector density, in this case only allow two RACE ports on a 6U VME board. Therefore part of the switching network is forced onto the backplane (a so called "active backplane") on which six switch chips reside. As systems expand to incorporate more processors to meet further mission requirements, active backplanes become even less attractive as they are very difficult to design, build, and maintain.




Figure 6.3-1: A 30-node system in 6U VME chassis.

To meet the challenging processing requirements of the near future, Mercury has prepared a "large system plan" that has 128 nodes per chassis, and four RACE ports per board to handle data between boards within a chassis. Board-to-backplane connector density is the limiting factor that force the use of large form factor "9U" (15x15 inch square) boards, with a twenty-slot active backplane in each chassis. Though preferred by system end users, a smaller "6U" chassis or a passive backplane is not currently possible, and could only be implemented in the future by using optical interconnects because of board-edge electrical connector density limitations. Optical interconnect technology will allow more processing power on each board and in each chassis, reducing the overall system size, and eliminating the active backplane (as shown in Figure 6.3-2).




Figure 6.3-2: Optical interconnects for IOP Mercury system

Optical interconnect technology for the IOP intra-chassis interconnects, or optical backplane offers three specific benefits:

- (1) It reduces the size of the system from 9U to 6U chassis - over 75% savings in volume.
- (2) It eliminates the use of an active electrical backplane, and replace it with a passive optical backplane.
- (3) It allows a more desirable system topology based on a central switched architecture.

6.3.2 The IOP STAP Processor

The IOP STAP processor configured in a 6U chassis is shown schematically in Figure 6.3-3. It consists up to 18 board slots for two corner turn (CT) boards, one to three FFT boards (determined by the availability of an FFT accelerator), eight slots for STAP processors, and up to five for Mercury off-the-shelf boards. All boards will be interconnected via a passive optical backplane where all the high speed, high density data will be carried by polymer optical waveguides, and low speed control signals and power can be carried by regular electrical interconnects. All off-the-board data I/O will be carried by "optical" RACE ports which will be described in the next section.




Figure 6.3-3: A schematic IOP processor in a 6U chassis with optical backplane.

6.3.3 The IOP backplane optical interconnect link design:

The new Mercury system which employs optical backplane interconnect technology will use a interconnect topology that requires a minimum of four RACE ports for each board-to-backplane connector or interface. For a system that is interconnected via one (or more) central switch board(s), the number of ports that is required for the switch board is likely to be much higher, i.e. a switch board that connects N boards will require $4N$ ports.

Figure 6.3-4 is a schematic illustration of a board with four optical RACE ports as the board-to-backplane interface. Compared with a standard electrical RACE based router chip, this board has a router with electrical RACE ports for intra-board interconnect (between processing elements on the board), and four optical RACE ports for off-board interconnect (between processing nodes through the backplane). The lower part of the figure shows a functional block diagram of the optical interface which translates a standard RACE port of 40-bit wide 40 Mbps electrical signals into a 3-bit wide 800 Mbps optical signal. This optical interface will replace the standard line drivers and receivers of an electrical RACE port.




Figure 6.3-4: A schematic illustration of a board, and a router with optical RACE ports.

The physical implementation of any optical link involves two key pieces of technology: (1) the optical transmitter and receiver (Tx/Rx) modules, and (2) the optical interconnect medium. In the context of an IOP optical backplane link, the first piece will be a Tx/Rx module based on VCSELs and integrated receivers packaged with a Si-CMOS router ASIC, or preferably, a new router module with an optical interface. The second piece will be a polymer optical waveguide-based backplane as the interconnect medium which includes a high-density board-to-backplane connector.

6.3.4 Optical RACE router module

We believe that a successful intra-cabinet optical interconnect solution for the IOP must take an approach which is based on low power dissipation and low cost design. Therefore, we have studied various choices of components, IC design options, packaging technologies, and their implications to performance, power, reliability, scalability, and cost of the whole link.

6.3.4.1 Power considerations

There is a minimum power required to run an optical link which is determined by the slope efficiency of the laser, the link loss budget, and the responsivity (gain) of the receiver. We estimate this power to be in the 10-30 mW per link range for link speed up to Gbps. Most of the optical links demonstrated to date show much higher power consumption (about 50-150 mW) for the following reasons.

- (a) Laser drivers are not optimized, or designed to take advantage of the superior characteristics of VCSELs which require very little pre-bias current (couple of mA or less today), and have very good temperature stability.
- (b) Most components (mux/demux chips, transimpedance amplifiers, and post-amplifier) are in single chip packages (SCP), and are designed to have general purpose I/O. This results in unnecessary use of low impedance line drivers which are often the most power hungry parts of these circuits.
- (c) Most of the link designs today are for single channel serial links where package type, size, and power constraints are different from those required for backplane type links.

An effective way to reduce the power consumption is integration - in both chip level (ASICs), and packaging level (MCMs). The idea is to bring the different parts of the link close together therefor eliminating the need for several low impedance line drivers between the parts.

Our approach is to make a router chip with an optical interface, which is illustrated in Figure 6.3-5. This is a module or a package that includes three different kinds of chips - the GaAs VCSEL chip and the GaAs Rx front-end (both in forms of linear arrays), and a Si-CMOS router chip with "optical IO buffers" for such functions as VCSEL drivers, receiver postamps, and mux/demux, etc.




Figure 6.3-5: A schematic illustration of a optical router MCM.

This approach has following advantages:

- (a) Most of the function blocks of the optical link such as mux/demux, code/decode, laser driver and postamps etc. (the “optical IO buffers”) are merged into a router ASIC in Si-CMOS, reducing the level of packaging, and eliminating redundant line drivers.
- (b) Since the minimum power required for driving an optical link does not scale with data rates up to 1 Gbps or more, we will take advantage of the recent developments in fast CMOS technology that bring its speed to 500 to 1000 Mbps range.
- (c) Although the “optical I/O buffers” add power, implementing this buffer in CMOS is probably the most power efficient approach. It also replace the standard line drivers that were normally provided.
- (d) This approach reduces the function and complexity of the GaAs chips to their bare minimums resulting in high yield and reliability, and lower cost.

6.3.4.2 Cost and other considerations

Adding functions like the “optical I/O buffer” into the router ASIC, and packaging the router with the lasers and receivers as a multi-chip module (MCM) tends to drive up the cost. However, both optoelectronic devices (the VCSELs and the GaAs receiver array chips) are relatively simple by themselves and can be tested at the wafer level and can be supplied as known good die (KGD). This greatly reduces the risk and cost associated with traditional MCMs. They can be handled and bonded into the package in the same way as those for electrical ICs. The inter-chip interconnects are limited and straight forward between the router and the optoelectronic devices, thus no complicated (or costly) MCM substrates are required. Therefore we do not expect a significant cost and reliability penalty by the MCM approach.

6.3.5 Polymer optical waveguide, and backplane connector

The optical interfaces to the Tx/Rx package will be provided by attaching a flexible optical waveguide ribbon, with a 45-degree terminated end facet, to the VCSEL or Rx chip (as shown in Figure 6.3-6). The optical assembly process can be added after the completion of the electrical portion.

We have demonstrated an alignment feature fabricated on the VCSEL chip at the wafer level to achieve passive-alignment of a waveguide ribbon to a VCSEL array with coupling losses less than 1dB.




Figure 6.3-6: A schematic illustration the VCSEL-to waveguide interface

6.3.5.1 Link loss budget

The physical implementation of the optical link for the IOP backplane is typically a point-to-point link which starts at the source VCSEL and ends on the photodetector of the receiver. Linking the two points are the optical media which consist of three waveguide segments, and four interfaces, two between the devices (VCSEL and Rx) and the waveguide and two board-to-backplane connectors. A link optical power budget is shown in Table 6.1. For a typical link, a source VCSEL power of 2 dBm can be expected. The losses at four interfaces and three waveguide segments add up to be about 13.6 dB, assuming up to 10 inches of total waveguide length in the link. This leaves -11.6 dBm of optical power available at the photodetector of the receiver, which is about 8.4 dB above a typical receiver sensitivity (for a bit-error-rate of 10^{-9} to 10^{-12} at bit rate up to 1Gbps).

TABLE 6.1. IOP BACKPLANE OPTICAL LINK LOSS BUDGET

		<i>Worst</i>	<i>Typical</i>	<i>Best</i>	<i>Unit</i>
Source	VCSEL power output	1.0	2.0	3.0	dBm
Losses					
Interface 1:	VCSEL-to-waveguide	2.0	1.5	0.5	dB
segment 1:	Module waveguide loss (1,2)	0.6	0.5	0.3	dB
Interface 2:	Board-to-backplane connector.	4.0	3.0	2.0	dB
segment 2:	Backplane waveguide loss (1,3)	5.1	3.7	2.0	dB
Interface 3:	Board-to-backplane connector.	4.0	3.0	2.0	dB
segment 3:	Module waveguide loss (1,2)	0.6	0.5	0.3	dB
Interface 4:	Waveguide-to-detector	2.0	1.5	0.5	dB
Receiver	Power available to Rx	-17.4	-11.6	-4.5	dBm
	Typical Rx sensitivity	-20	-20	-20	dBm
Margin		2.7	8.4	15.5	dB
(1):	Typical waveguide loss	0.25	0.18	0.1	dB/cm
(2):	typical length less than 1 inch				
(3):	Length of 8 inches assumed				

Among the possible polymer waveguide choices available to the IOP program are Honeywell's polymer waveguides which are based on proven commercial materials (ULTEM and BCB) and standard electronics fabrication processes. The waveguides have survived, without performance degradation, standard board lamination process and Military tests (MILP139-49/13). The loss of this waveguide is 0.24 dB/cm. Dupont's Polyguide, which is also licensed to AMP, has a typical

loss of 0.2 dB/cm, and Allied Signal has reported a loss of their Acrylate monomer based waveguides to be 0.1 dB/cm or less. Since the losses of the highly multi-mode waveguides can be affected by many design, fabrication, and test conditions, there is an ongoing effort at Sandia National Lab to systematically evaluate these waveguides.

6.3.5.2 Bit error rate (BER)

Bit error rate (BER) of an optical link is fundamentally determined by the signal to noise ratio (SNR) at the receiver. The signal level arriving at the receiver input is limited by laser output power, and the link loss budget. The noise of the link is typically dominated by the receiver front-end, and varies with data rate (bandwidth), and depends on detector and amplifier technology. For multimode links, there is potential modal noise. For parallel link, there is additional noise related to signal cross-talk. The noise associated with optical signals themselves is normally not data rate dependent.

There is an often referenced optical link parameter called "receiver sensitivity", It is defined as the minimum detectable optical signal to achieve certain BER at a given data rate.

Therefore a BER is function of signal to noise ratio (SNR), which is determined by Tx output power (STX), link loss (Link), and link noise (N) which is dependent on link data rate.

$$\text{BER} = f(\text{SNR}), \quad \text{where } \text{SNR} = (\text{STX} - \text{Link}) / N$$

Statistics show that, under normal circumstance, BER(Q) is 10^{-9} , 10^{-12} , and 10^{-15} when Q is approximately 6, 7.04, and 7.94, respectively, where Q is equivalent SNR of optical power of 7.8, 8.5, and 9 in dB. We notice that the BER is very sensitive to SNR. A 1.2 dB SNR improvement can change the BER from 10^{-9} to 10^{-15} , or vice versa. A 10^{-15} BER for a link operating at 800 Mbps means there will be one error every 14.5 days of continuous operation. In the typical case of the IOP link budget where we have over 8 dB of margin for the link, an error-free link can be expected.

6.3.5.3 Board-to-backplane connector, and its density

A waveguide based high density backplane connector is currently in development under a DARPA funded program POINT, which involves Honeywell, AMP, GE, Allied Signal, UCSD and Columbia University. The objective of the POINT program is to develop key building blocks, such as a practical board-to-backplane connector, for a high density optical backplane. Our goal is to demonstrate an optical board-to-backplane connector that offers 5-10 times the density of today's state-of-the-art electrical connector.

To date, we have demonstrated a waveguide connector for 144 optical channels at a channel pitch of 100 μm (see figure 6.2-2). The physical width of this connector is less than 1.5 inch. Considering our optical RACE port approach requires six optical channels, we will be able to fit equivalent of 24 optical RACE ports through this compact connector. If we implement the backplane connector on a centralized switch board where large number of interconnect ports are

needed at board edge, the connector is enough to interconnect with six other boards with four RACE ports per each board. If we utilized half of the 9-inch board edge of a 6U board, we could accommodate at least three such connectors which will interconnect all $6 \times 3 = 18$ boards that can be put into a standard chassis. Of course, the real interconnect capacity of such a proposed central switch board will most likely be limited by power or space available of fitting large number of router chips on to the 6U board.

6.4 IOP OPTICAL BACKPLANE INTERCONNECT RELIABILITY ISSUES

In the past few years, there have been several parallel optical interconnect module prototypes and link demonstrations, such as **OETC**, **POLO**, **OPTOBUS**, and **POINT**, made by various companies including Honeywell. However, reliability data neither for the parallel modules nor for the links are available at the date of this report since no company is currently producing such parallel links as products. While the complete optical link reliability is a complex issue related to the various components, packaging approaches, and operating conditions and environments, we will comment in this section on the reliability of components including VCSELs, receivers, optical router, optical waveguides and connectors.

6.4.1 VCSELs

VCSELs have become the key optoelectronic component for Honeywell's optical interconnect technology. Honeywell is the first company that has introduced VCSELs as qualified (with complete reliability tests) commercial product for optical data communication applications. VCSELs are now in volume production (in 100,000s) in Honeywell's Micro Switch Division. They are being designed into their next generation high speed data communication modules, as well as being designed into their next generation high speed data communication modules as well as being supplied as packaged components to other users. Current production is focused on discrete single element VCSEL devices for serial data link applications. However, the nature of the device allows one to easily scale the product to 1D or 2D array devices with a simple layout change. High yields of the devices (99.8% and 94% for discrete and 1×32 array devices, respectively, have been achieved, and easy wafer level testing keeps the cost of this component competitive.

Honeywell MICRO SWITCH has a long history of providing highly reliable, superior quality products. Honeywell takes four basic approaches to ensure high reliability of its products. First, its development program for new products includes extensive reliability simulation and analysis. Second, the development program includes exhaustive reliability testing, such as sensitivity to electrostatic discharge (ESD), mechanical and thermal tests, and accelerated life testing. Third, a sample from each wafer is subjected to reliability testing, including burn-in, before release of the wafer to production. Also, production processing includes environmental stress screening, typically temperature cycling and burn-in, for products as necessary to ensure good reliability for devices shipped. Fourth, Honeywell continues to monitor product reliability and supplements the reliability database through its Product Reliability Monitor program, which periodically subjects a large sample of production devices to a battery of reliability tests, including extended

operating life testing.

In the paragraphs that follow, we briefly describe the scope of the reliability tests that have been done and are continuing at Honeywell.

Electrostatic Discharge Study. The electrostatic discharge (ESD) sensitivity of a product is of considerable importance, especially since devices with small active regions, like the VCSEL, are typically susceptible to ESD-induced damage. For these reasons, human body model ESD testing was performed on VCSELs from multiple wafers using MIL-STD-883D, method 3015.7. Subsequent 2000-hour burn-in was done to assess any latent defect caused by the ESD shock. The failure threshold was found to be above 700 V, which is somewhat more robust than identical testing showed for edge-emitting semiconductor lasers. It is important to note that no significant latent ESD effect was found during the burn-in.

Mechanical Testing. Mechanical testing has included air-to-air temperature cycling (100 cycles, -40°C to 100°C, 5 second transition time) per MIL-STD-750, method 1051 for 1051 for 120 HFE4080 parts with no failures. Additionally, the same parts passed after constant acceleration testing (20kg, Y1 axis) per MIL-STD-750, method 2006. Mechanical testing continues as part of our product reliability monitoring program.

Life Testing Summary. The principal reliability data is from over 500 metal TO-packaged VCSEL devices subjected to operating life testing in 22 burn-in groups at five temperatures and five operating currents. This study comprised eight wafers built during a span of more than a year, compiling nearly two million actual burn-in-device-hours. These figures do not include additional reliability testing done on other VCSEL parts, which takes the total number of device-hours to nearly 3.5 million as of March 1997 (not including production burn-in). The failure criterion selected was a 2 dB drop (or 1 dB increase) in total optical output power relative to the power measured before starting life testing. Analysis of the data showed that a 3 dB criterion would give about 10-20% longer lifetimes than the 2 dB criterion. An extremely important point is that all failures recorded were actual failures and not extrapolations to failure.

Statistical data analysis shows that the device reliability, described in mean time to failure (MTTF) is a function of VCSEL driving current and junction temperature (which in turn is a function of ambient temperature, driving current, and data rate). For example, for a typical 10 mA of driving current and 40 degrees C junction temperature, the MTTF is 2.8×10^7 hours, or approximately 3,200 years. This is equivalent to 35 FIT, ($1 \text{ FIT} = 10^9 / \text{MTTF}(\text{hr.})$). This figure is shown in Table 6.4.1 where we summarize the reliability data collected for various components. Additional reliability testing is continuing, and confirms that on-going device enhancements are further improving the reliability of Honeywell's VCSEL.

TABLE 6.4.1. SUMMARY OF OPTICAL LINK COMPONENTS RELIABILITY

Device	Technology	Reliability	Remarks
VCSELs	GaAs - MOCVD	2.8x10 ⁷ Hrs. MTTF or 35 FIT	Honeywell Micro Switch (10 mA, T _j =40C)
Integrated receiver	GaAs - MESFET	40 FIT 7 FIT	Triquint Semiconductors (150 C) ⁽¹⁾ Vitesse Semiconductors (100 C) ⁽¹⁾
Router ASIC	Si - CMOS	59 FIT	Cypress Semiconductors (150 C) ⁽¹⁾ Current supplier of Mercury Router)
Polymer waveguide	ULTEM/BCB on Kapton	survives MilP-139-49/13 test	Rugged
Optical backplane connector	MT like or Super-MT	TBD	AMP, 3M as potential suppliers
Router package	MCM	TBD	VCSEL and photodetector can be packaged using standard processes

(1): Taken from manufacturers literature of similar products.

6.4.2 Integrated Receiver

The integrated optical receiver arrays used in our optical link design are now available from commercial GaAs MESFET foundries, such as Vitesse Semiconductor and Triquint Semiconductor. This is a device with an MSM photodetector integrated with E/D MESFET amplifier circuits. Incorporating photodetectors into the IC requires little or no modification to the existing MESFET processes, therefore the reliability of such an integrated receiver chip should be comparable with that of their catalog products. In Table 6.4.1, we list some of the reliability data quoted in these manufacturers' digital IC product catalogs.

It should be noted that GaAs ASIC technology has been growing at a double-digit rate in recent years. The two companies mentioned above alone have combined annual revenue of about \$200 million in 1996. The GaAs IC technology has become one of the mainstream semiconductor technologies for RF, wireless, and digital communication applications.

6.4.3 Optical Router Chip

The router chip with optical interface described in our design is a standard crossbar switch chip with its IO buffers tailored to interface with optoelectronics devices. The whole IC is intended to be implemented as an ASIC using standard Si CMOS technology for its low power, low cost, and good reliability characteristics. The six-port router ASIC that Mercury Computing Systems uses in their products today is a COTS chip fabricated by Cypress Semiconductor. As a reference in Table I, we quote reliability data of a product from Cypress' catalog that is similar to Mercury's router.

6.4.4 Optical polymer waveguides and backplane connectors

Other key components for optical backplane interconnects include polymer optical waveguides and backplane connectors. These devices have been developed under several government funded demonstration and prototype programs in recent years. Currently, several commercial companies such as Honeywell, GE, Allied Signal, AMP, and 3M are actively pursuing their productization.

Honeywell's Ultem/BCB waveguides are based on materials that have proven stability and reliability over wide temperature range, and are widely used in microelectronics and other industries. They are typically made on substrates like Kapton which is a standard flexible circuit board material. There are also several other polymer technologies such as Dupont's Polyguide, and Allied Signal's Acrylate monomer based waveguides which are also very promising for optical backplane interconnects. Packaging technology is another key element of the overall link reliability. The optical router module proposed in the IOP Program is an assembly of optoelectronic chips, optical waveguides, and backplane connectors. We point out that the electrical interconnect pad configurations of the optoelectronic devices such as VCSELs and integrated receiver chips are identical to that of the electrical ICs. They can be packaged into a module using the same processes and techniques used for any electronic ICs, such as wire, tap, or flip-chip bonding. We have also developed optical interfacing approaches based on passive alignment techniques for low cost and reliable assembly. Optical waveguide based backplane connectors using expanded-beam approach allow misalignment tolerances that are compatible to existing electrical backplane connectors.

6.5 SUMMARY

The reliability data available today show that VCSELs and GaAs receivers are of better or similar reliability characteristics than that of today's electronic ICs. While reliability data of polymer waveguides and connectors are not available today, our optical interconnect approach, based on a design philosophy that uses the same proven materials, processes, and assembly techniques of the electronics systems, assures practical, reliable results.

We have identified the optical interconnect technology and the baseline approach for the IOP STAP processor. Our approach will reduce the system size by 75% and eliminate the need for active backplanes. The high density optical backplane will also allow a central switched system architecture. High density optical board-to-backplane connectors will allow the equivalent of up to 72 or more RACE ports over 4.5 inches of linear board edge. This technology gives system designers the flexibility to build network topologies that are not possible today.

7. ACCELERATOR NODES

High throughput-density is fundamental to the realization of a flyable real-time STAP processor. The size, weight, and power limitations inherent in an airborne environment, as well the data communications problems of standard COTS-based approaches to STAP identified earlier in this report demand a more efficient STAP solution.

Our IOP design relies strongly upon a set of highly efficient STAP accelerators to deliver the orders of magnitude improvements in processor complexity needed to make airborne STAP viable. Residue Number System (RNS) processing techniques, with their promise of high speed and efficient handling of complex arithmetic appear strongly suited to the matrix operations required as part of the STAP algorithms.

Our analysis of the IOP's STAP requirements has led to the selection of a set of four STAP accelerator modules. This set of candidates for hardware acceleration have been carefully chosen so as to maximize their general applicability. By identifying and accelerating a set of general STAP "building blocks", rather than a specific STAP algorithm itself, our chosen accelerators are useful in a wide variety of STAP applications (e.g. they are capable of performing all three Mountain Top STAP algorithms). They are parametrically programmable (number of receiver channels, degrees of freedom, range cells, etc.), and are modularly constructed, so that they may be cascaded in either dimension to solve larger STAP problems.

Our proposed accelerator modules include:

- Corner-Turn: Performs the matrix transpose required as part of Doppler processing. Eliminates the need to perform the distributed corner turn, found to be the major data flow bottleneck in COTS-based STAP. Using off-the-shelf components, the corner turn board occupies a single 6U VME card.
- FFT Module: Performs the Fourier Transform required as part of Doppler processing. This module employs off-the-shelf components (e.g. Sharp FFT chips). (A RACEway compatible FFT accelerator may become available off-the-shelf in the EOA development timeframe, eliminating the need for this design.)
- Weight Maker: Calculates adaptive weights via Cholesky Decomposition. Employing two RNS processing ASICS, this module executes the equivalent of more than 140 Gigaops on a single 6U VME module (the desire to employ air-cooling will likely require partitioning the design across several boards as described later in this section.)
- Beamformer: Applies the raw receiver data. Employs the same ASIC devices as the Weight Maker.

System Overview:

The accelerator nodes consist of the FFT module, the Weight Maker Module, and the Digital Beamformer Module. A preprocessing node is located near the A/D converters which performs FIR filtering and discrete Hilbert transforms. Two Corner Turn Memories are also inserted before and after the FFT Module. Optical interconnects are utilized between the accelerator nodes (see Figure 7.0-1).

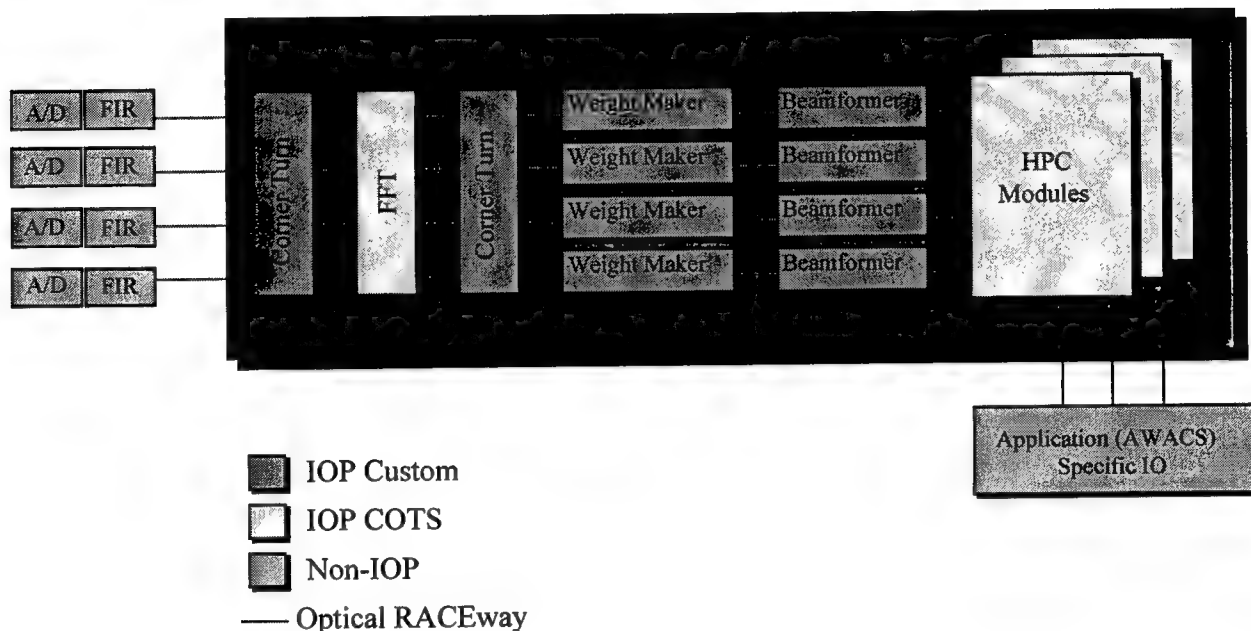


Figure 7.0-1 IOP System Block Diagram

There are 128 A/D converters that produce digital 16 bit real data at 50 MHz. The preprocessing node changes this to 128 channels of 20 bit complex data at 1 MHz (5.1 Gbps). The optical interconnect between the preprocessing node and the first Corner Turn Memory changes the data to 8 channels of 20 bit data at 32 MHz (5.1 Gbps).

From the first Corner Turn Memory to the FFT Module the data is 4 channels at 40 bits and 32 MHz (5.1 Gbps). From the FFT to the second Corner Turn Memory there are 8 channels of 24 bit data at 32 MHz (6.1 Gbps), and the data from the second Corner Turn Memory to the Weight Maker Module and the Beamformer Module is the same. The data rates from the Weight Maker Module to the Beamformer Module and from the Beamformer Module to the COTS HPC Modules are relatively small, even though there are 8 simultaneous beams. Figure 7.0-2 shows a summary of the system data rates.




Figure 7.0-2 IOP System Data Rates

The CPI of the IOP system is 200 ms with 8 ms of space charge. The number of degrees of freedom is 128 and the range cell rate is 1 MHz. This means that there are 192,000 data vectors in each data cube. There are three Doppler modes and the number of Doppler bins in the three modes are respectively 64, 128, and 256, and the number of range cells in the three Doppler modes are 3000, 1500, and 750. Figure 7.0-3 shows the data cube for the first Doppler mode.




Figure 7.0-3 STAP Data Cube

The number of data matrices per Doppler bin (must have integer value) in the three Doppler modes are respectively 6, 3, and 2. The number of samples per data matrix in the three Doppler modes are respectively 500, 500, and 375 and the number of data matrices per data cube in the three Doppler modes are 384, 384, and 512.

The third Doppler mode is the oddball because we would want the number of data matrices per Doppler bin to be 1.5, but are forced to use the integer value of 2. This makes the third Doppler mode more difficult than the first two. In the final design, we may choose to process a single

data matrix per Doppler bin for the third mode. This would make the number of data matrices 356 and the number of samples per data matrix 750. The third Doppler mode then becomes easier than the first two.

7.1 DESIGN TRADES

The most important consideration in the IOP system is the number of operations required by the Weight Maker Module. In the first two Doppler modes, a 128 by 500 complex data matrix must be processed every $200 \text{ ms}/384 = 521 \text{ } \mu\text{s}$. A binary algorithm based on orthogonal transformations would require at least $2 \times 500 \times 128^2$ complex multiply-adds to process one data matrix. If we count an operation as a real multiply or a real add, then each complex multiply-add equals 8 operations. Therefore to process a data matrix every $521 \text{ } \mu\text{s}$ equals a processing rate of more than 250 GOPs. Figure 7.1-1 shows the system operation rates.

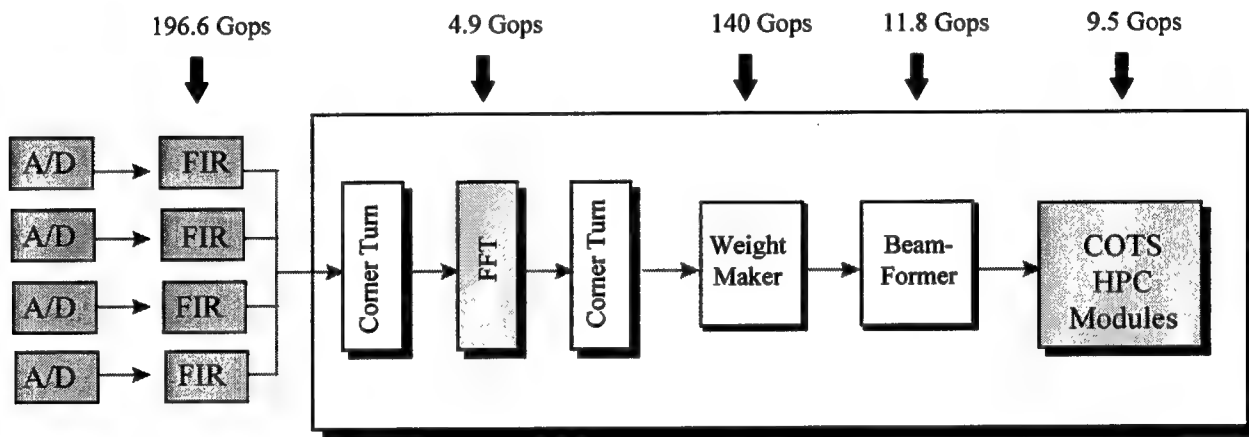


Figure 7.1-1 IOP Throughput Requirements

In a conventional COTS-based STAP processor this would amount to thousands of processing elements, so a more highly optimized solution for the Weight Maker Module is required. The solution requires that multiple processing elements be placed on a single ASIC and that data sharing among ASICs be minimized.

One essential simplification is to allow the Weight Maker Module to process one data matrix at a time instead of one data cube at a time. The Corner Turn Memory inserted after the FFT module makes this possible. It makes sense to also insert another Corner Turn Memory before the FFT module since the two are nearly identical.

The proposed Weight Maker Module incorporates a unique synthesis of arithmetic, algorithm, and implementation. Two custom ASICs are proposed for this module each with about 400K useable gates. The entire Weight Maker Module can fit on a single 6U board when the clock rate for the ASIC is 100 MHz, however due to VME chassis cooling restrictions, it will likely be partitioned across multiple 6U boards.

7.1.1 RNS Arithmetic

The residue number system (RNS) is an integer number system that allows fast parallel implementation of arithmetic over long wordlengths. The factored look-up table technique for RNS arithmetic was first developed by Westinghouse in the Optical Adaptive Program (OAP) and refined by SRC for an electronic implementation. In this technique the moduli that can be used must be a prime of the form $p = a*b+1$ where a and b are relatively prime and are each less than or equal to 16. Furthermore if the modulus p equals 1 modulo 4 then the quadratic residue number system (QRNS) can be employed.

Figure 7.1.1-1 shows the circuit for multiplication and addition for a modulus p that has the form $p=a*b+1$. The modulus set used by the Weight Maker Module is: {241, 157, 113, 89, 73, 61, 53, 41, 37, 29, 17, 13}. This gives the equivalent of about 70 bits of binary dynamic range but the actual wordlength of a number in RNS is 75 bits. The moduli used in the Weight Maker Module have the following factorization:

$$241 = 16 \times 15 + 1$$

$$157 = 12 \times 13 + 1$$

$$113 = 16 \times 7 + 1$$

$$89 = 8 \times 11 + 1$$

$$73 = 8 \times 9 + 1$$

$$61 = 4 \times 15 + 1$$

$$53 = 4 \times 13 + 1$$

$$41 = 8 \times 5 + 1$$

$$37 = 4 \times 9 + 1$$

$$29 = 4 \times 7 + 1$$

$$17 = 16 \times 1 + 1$$

$$13 = 4 \times 3 + 1$$

Note that each modulus is 8 bits or less and each factor is 4 bits or less. This means that all logic functions associated with this arithmetic have 8 or fewer input bits. The logic functions associated with this RNS arithmetic have been synthesized and the gate count for a multiplier and adder is about 6K. An equivalent 35 bit binary multiplier and adder would require at least 12K gates.

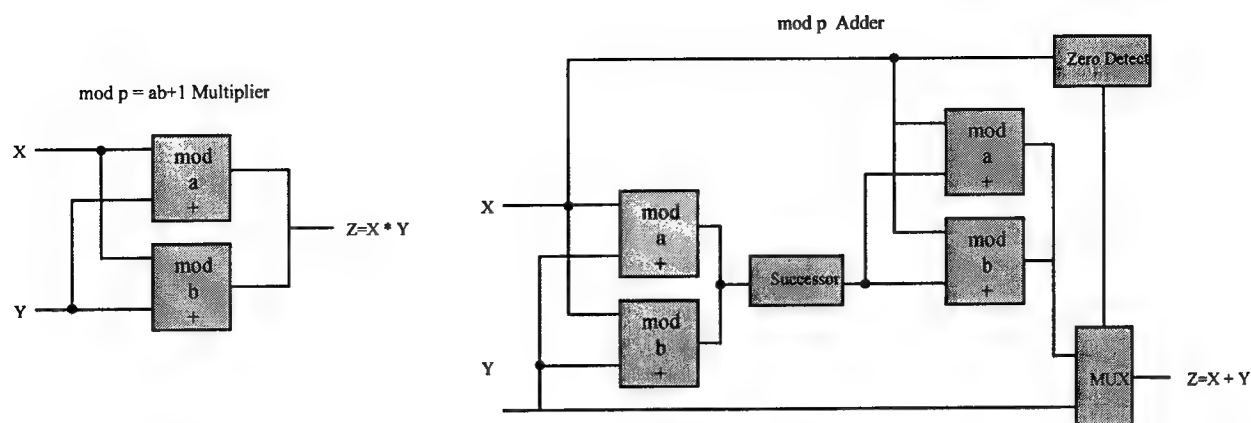


Figure 7.1.1-1 Factored RNS Multiplier / Adder Schematic

The RNS is an integer number system in which only the exact integer operation of multiplication and addition are superior to binary. Other operations necessary for STAP processing include scaling, threshold detection, and inverse square root. Efficient implementations of these operations and for RNS to binary conversion have been developed by SRC. These implementations are based on core functions.

7.1.2 RNS Choleski Algorithm

The RNS Choleski algorithm was specifically designed to solve the STAP problem in RNS. The main feature of this algorithm is that the complexity of the core-based operations in RNS is only quadratic in the number of degrees of freedom whereas the complexity of the exact integer operations is cubic. This means that the vast majority of the operations are the exact integer operations for which RNS is best suited.

The operation count of the RNS Choleski algorithm is one fourth that of a comparable binary algorithm. A factor of two comes from the fact that the covariance matrix is formed instead of using orthogonal transformations such as Householder or Givens. This is actually a mixed blessing because a longer wordlength is required since the condition number of the covariance matrix is the square of the condition number of the data matrix. An additional factor of two comes from the utilization of QRNS which reduces the number of operations in each complex multiply-add by a factor of two.

The RNS Choleski algorithm with a 70 bit RNS dynamic range is comparable in numeric precision to a 35 bit binary algorithm using orthogonal transformations. The 70 bit RNS multiply-adder needs less than half the gates of a 35 bit binary multiply-adder.

In binary it would be impractical to switch to a Choleski type algorithm because doubling the wordlength would quadruple the cost of each multiplier and double the cost of each adder. Also there is no counterpart to the QRNS in binary. In RNS it is not possible to switch to an algorithm based on orthogonal transformations because too many core-based operations would be required.

Figure 7.1.2-1 shows the functional diagram of the RNS Choleski algorithm. A 128 by 500 complex 24 bit binary data matrix is first converted to an RNS 128 by 500 complex 75 bit data matrix. The complex RNS data matrix when converted to QRNS becomes a pair of real matrices D1 and D2, where D1 is 128 by 500 (75 bits) and D2 is 500 by 128 (75 bits). In the QRNS domain, the covariance matrix R is the product of D1 and D2.




Figure 7.1.2.1 Choleski RNS Weight Calculation Algorithm

In RNS and in binary, the covariance matrix is a 128 by 128 conjugate symmetric complex matrix, but in QRNS it is a 128 by 128 nonsymmetric real matrix. Making the covariance matrix requires the majority of the operations in the RNS Choleski algorithm, and these are all exact integer operations.

The next step of the algorithm is the Choleski decomposition. In this step the covariance matrix is overwritten with the Choleski factor. In binary and in RNS, the Choleski factor is a complex lower triangular 128 by 128 matrix. In QRNS it is a real 128 by 128 matrix. The core-based operations required are threshold detection and inverse square root on the diagonal elements, and two scalings for the off diagonal elements. The number of exact multiply-adds in this step is still cubic with the number of degrees of freedom, but only one sixth as many as are needed to make the covariance matrix. The final step in the algorithm is the forward and backward elimination which are performed for eight different steering vectors. The only core-based operations required are scalings, and the number of them is only linear with the degrees of freedom. Exact multiply-adds have a quadratic complexity with the number of degrees of freedom in this step.

7.1.3 RNS STAP Implementation

An efficient implementation of the RNS Choleski algorithm requires multiple processing elements on a single ASIC that can be simultaneously kept busy within reasonable I/O requirements. The formation of the covariance matrix is the most important task, and involves the multiplication of the real matrix D1 (128 by 500) with D2 (500 by 128) to form R (128 by 128).

This task involves $2 \cdot 500 \cdot 128 \cdot 128$ operations but only $2 \cdot 128 \cdot 500$ inputs, so the ratio of operations to inputs is 128. A high ratio of operations to inputs is precisely what is needed for

the efficient utilization of an ASIC with multiple processing elements. This is why the key to the RNS STAP implementation is the Matrix Multiplication ASIC.

At 6K gates per multiply-adder it was determined that the maximum number of multiply-adders that can fit on a single ASIC is 32. This means a single Matrix Multiplication ASIC can compute the 32 by 32 product of any two matrices of sizes 32 by N and N by 32 for any integer N. The Matrix Multiplication ASIC can keep all 32 multiply-adders busy and have an operation to input ratio of 32.

Sixteen Matrix Multiplication ASICs can be arranged in a 4 by 4 array to compute the entire 128 by 128 covariance matrix. Each ASIC would hold a 32 by 32 portion of the result. The covariance matrix is then processed in place in order to avoid excessive transfer of data. Only the data transfers necessary to overwrite the covariance matrix with the Choleski Factor and compute the weight vectors are made.

The Matrix Multiplication ASIC is packed with multiply-adders and RAM in order to efficiently perform the large majority of the operations in the RNS Choleski algorithm. This means that a second ASIC design is needed to perform the rest of the operations. The Core ASIC performs all the core-based operations and transformations. There are far fewer of these operations in the RNS Choleski algorithm than exact multiply-adds but they are necessary.

Early in the program, consideration was given to making the Matrix Multiplication ASIC operate over a single modulus. This would allow more multiply-adders to be placed on a single ASIC. A programmable modulus concept was developed to avoid a different ASIC design for each modulus. It turned out that the programmable modulus was not nearly as efficient in terms of gates as a custom design. Since multiple ASIC designs would be too expensive, the alternative was to put all the moduli on one ASIC and use custom logic for each modulus.

7.2 CORNER TURN MEMORY

The data coming from the preprocessing node is by element space in parallel, then by range, followed by pulse/Doppler. The FFT processing node wants the data by element space then by Doppler followed by range. The first Corner Turn Memory permutes the data and introduces one data cube or 200 ms of latency to the system. The Weight Maker Module wants the data by element space, then by range, followed by Doppler. The second Corner Turn Memory again permutes the data and introduces another 200 ms of latency.

Figure 7.2-1 shows a schematic of the Corner Turn Memory. This memory uses COTS components and will fit on a single 6U card. The size of this memory is 2 Gbits. The board I/O count is 384 data bits consisting of eight 24-bit input buses and eight 24-bit output buses, each running at 32 MHz, plus control and some clocking.

7.3 FFT MODULE

Because of the first Corner Turn Memory, the FFT Module may consist entirely of dedicated FFT chips that output the FFT'd data in the same order as the input. The alternative is to use DSP chips that can both permute the data and perform the FFTs. This alternative might be considered if it can eliminate the 200 ms latency that the first Corner Turn Memory introduces.

Depending on the Doppler mode, the FFT module must compute 384,000 length 64 FFTs, or 192,000 length 128 FFTs, or 96,000 length 256 FFTs every 200 ms. The processing rates are respectively 3.7, 4.5, and 4.9 GOPs. A practical solution is to use COTS Sharp FFT chips. The current 40 MHz Sharp FFT chip (LH9124) computes a 256 point complex FFT in 16.2 μ s. The next generation 100 MHz Sharp FFT chip will be available very soon.

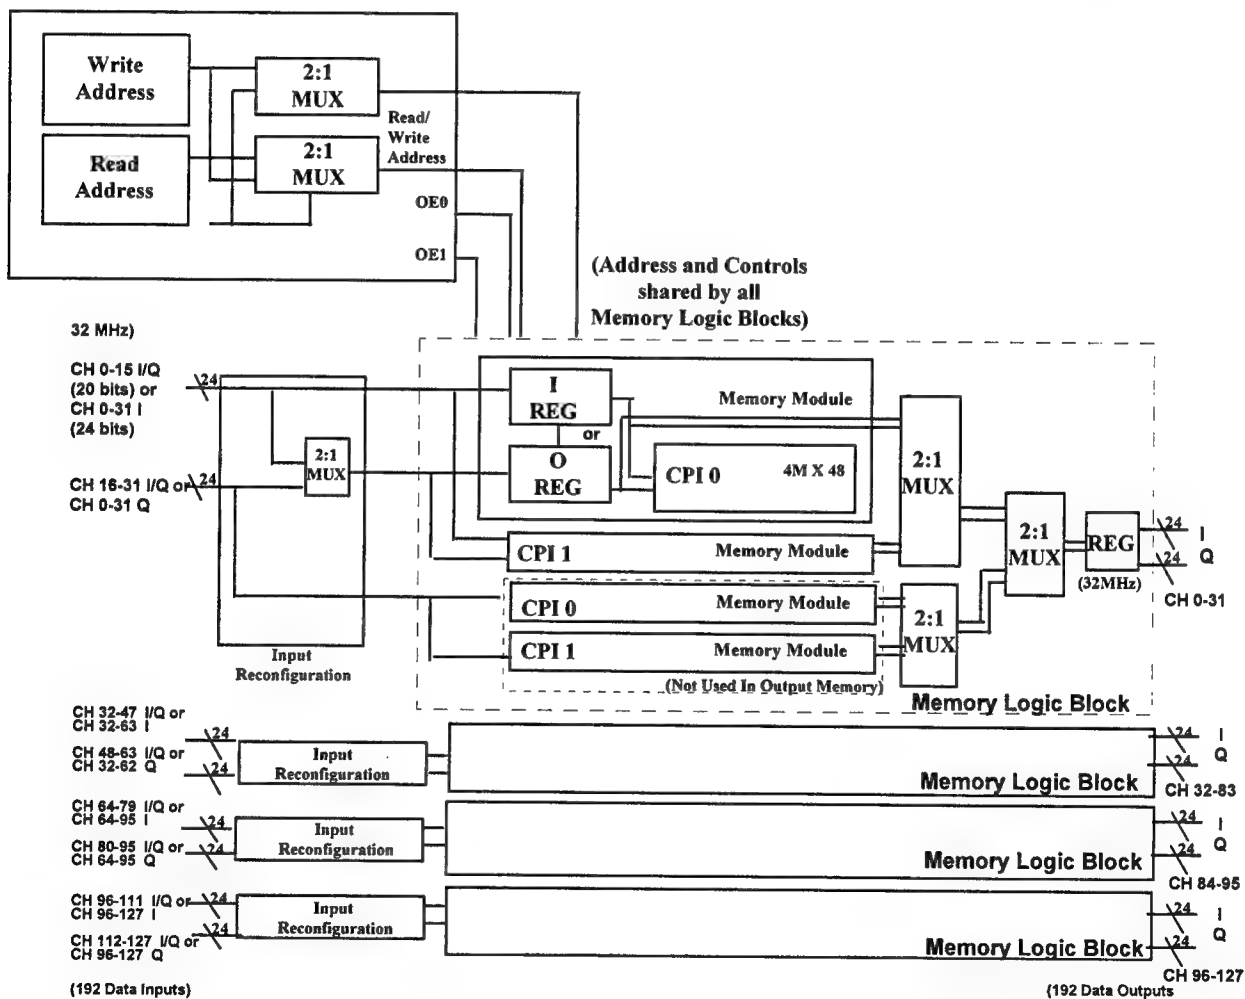


Figure 7.2-1 Corner Turn Memory Module

7.4 WEIGHT MAKER MODULE

The Weight Maker module must process a 128 by 500 data matrix every 521 μ s. The proposed module for the IOP uses two custom RNS ASICs. Figure 7.4-1 shows a schematic of the Weight Maker Module. It consists of sixteen Matrix Multiplication ASICs, four Core ASICs, four 32K by 48 bit memories (6Mb of RAM), and some control and glue logic.

The data that goes into the Weight Maker Module first enters into the RAM. The RAM is a ping pong memory that holds two data matrices so that one data matrix can be loaded in the time required to process the other.

The data matrix is fed to the Core ASICs which perform binary to RNS and RNS to QRNS conversion and feed the QRNS data to the 4x4 array of Matrix Multiplication ASICs. This process will take 160 μ s during which time the covariance matrix is formed within the Matrix Multiplication ASICs.

The Choleski decomposition is performed next. This involves calculations by both the Matrix Multiplication ASICs and the Core ASICs and some data transfers between them. The Core ASICs must perform scaling, threshold detection, and inverse square root operations. This process takes about 130 μ s. Finally, a forward and backward elimination process is performed to compute the eight weight vectors. The forward and backward eliminations take about 210 μ s.

The next data matrix is then ready to be processed in about 500 μ s whereas 520 μ s is available. By far the most efficient portion of the algorithm is the formation of the covariance matrix in terms of the number of operations performed in the time consumed. Next is the Choleski decomposition and the least efficient is the forward and backward elimination.

In the third Doppler mode with the number of data matrices per Doppler bin equal to 2, there would be only 391 μ s for each data matrix. The 160 μ s required to make the covariance matrix would drop to 120 μ s because the number of samples drops from 500 to 375. The time required for Choleski decomposition and forward and backward eliminations would remain the same. The Weight Maker Module would require 460 μ s but would only have 391 μ s.

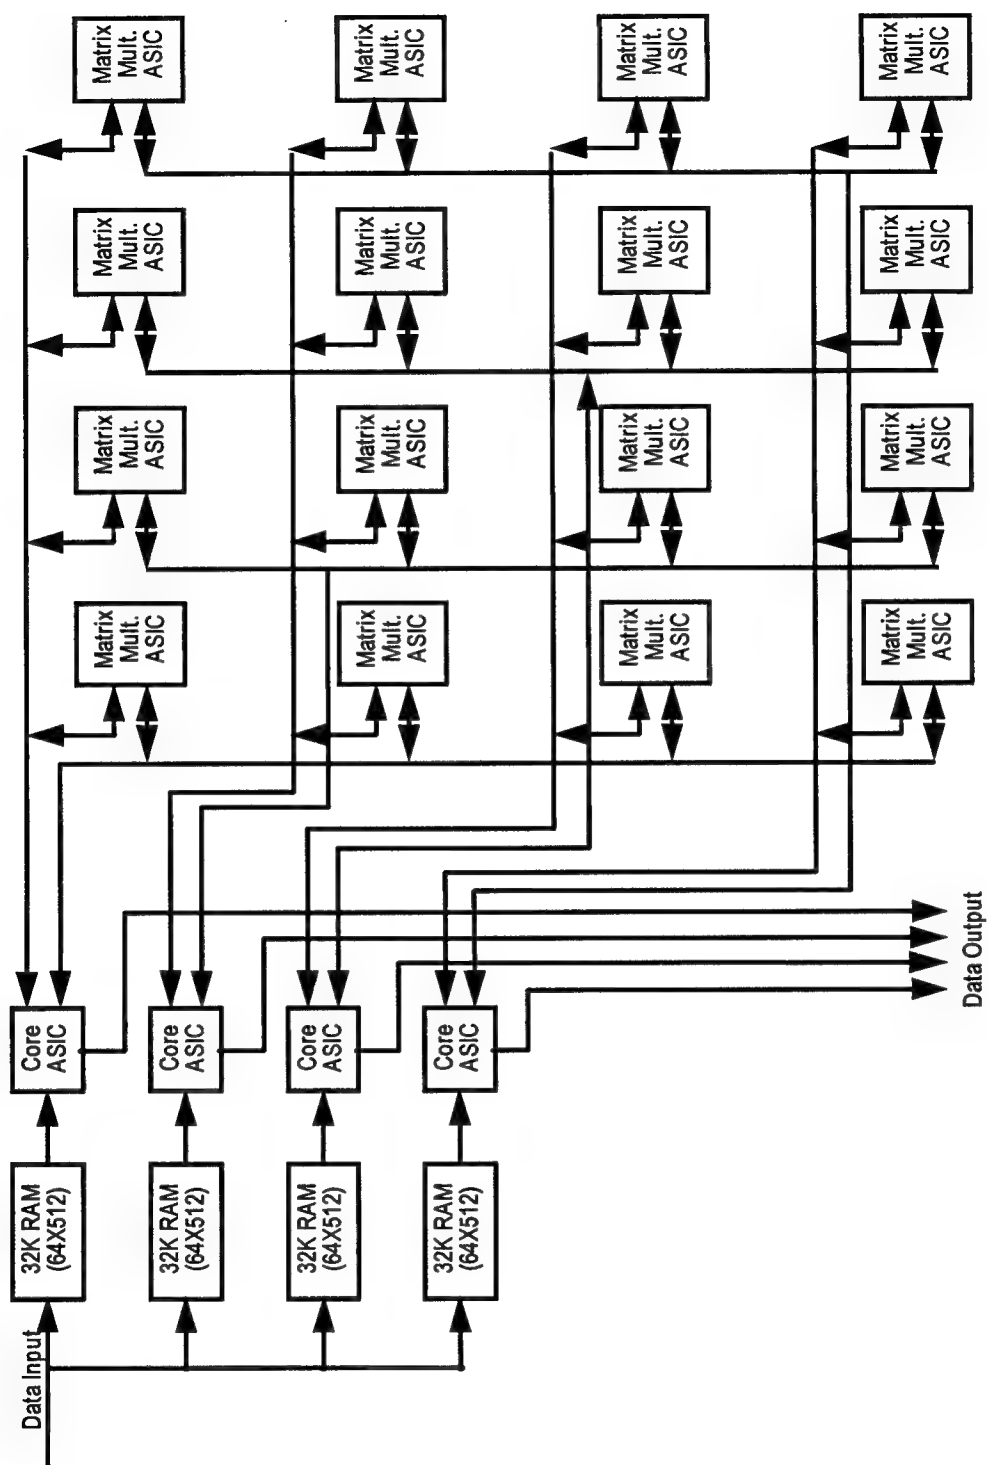


Figure 7.4-1 Weight Maker Module

A solution to make the third Doppler mode work, is to reduce the number of data matrices per Doppler bin to 1. Then 781 μ s would be allowed per data matrix and the time required to form the covariance matrix would increase from 160 μ s to 240 μ s. This means that the Weight Maker Module could process a data matrix in 580 μ s.

The Weight Maker Module can easily handle any number of degrees of freedom up to 128 and any number of samples per data matrix. The time required to make the covariance matrix varies quadratically with the number of degrees of freedom and linearly with the number of samples per data matrix. The time required for the rest of the algorithm varies cubically with the number of degrees of freedom. The most efficient part of the process is the making of the covariance matrix so adding more samples per data matrix adds very little cost.

To process more degrees of freedom would require a larger board, or multiple boards with interconnections. The current design is limited to a maximum of 128 degrees of freedom.

7.4.1 Matrix Multiplication ASIC

The matrix multiplication ASIC is the workhorse of the Weight Maker Module. Figure 7.4.1-1 shows a schematic of this ASIC. The ASIC contains 32 multipliers and adders and 32 RAMs. Each multiply-adder computes one row of a 32 by 32 matrix product and stores it in the RAM. The Matrix Multiplication ASIC requires only two input words per clock but keeps all 32 multiply-adders busy computing a 32 by 32 portion of the covariance matrix that is stored in the distributed internal RAM. With 16 Matrix Multiplication ASICs operating at the same time, the covariance matrix is computed in 160 μ s, the same time required to load a data matrix.

7.4.2 Core ASIC

The Core ASIC is responsible for the core-based operations scaling, threshold detection, inverse square root, and RNS to binary conversion. The Core ASIC is also responsible for the transformations from binary to RNS, RNS to QRNS, and QRNS to RNS as well as some exact multiply-adds. The operation rate of the Core ASIC is only twice the clock rate of the ASIC which is why it is important that the algorithm minimize the core-based operations.

A schematic of the Core ASIC is shown in Figure 7.4.2-1. The inverse square root function is shown in Figure 7.4.2-2, binary to RNS conversion in Figure 7.4.2-3, RNS to binary conversion in figure 7.4.2-4. RNS to QRNS and QRNS to RNS conversions are shown in Figure 7.4.2-5.

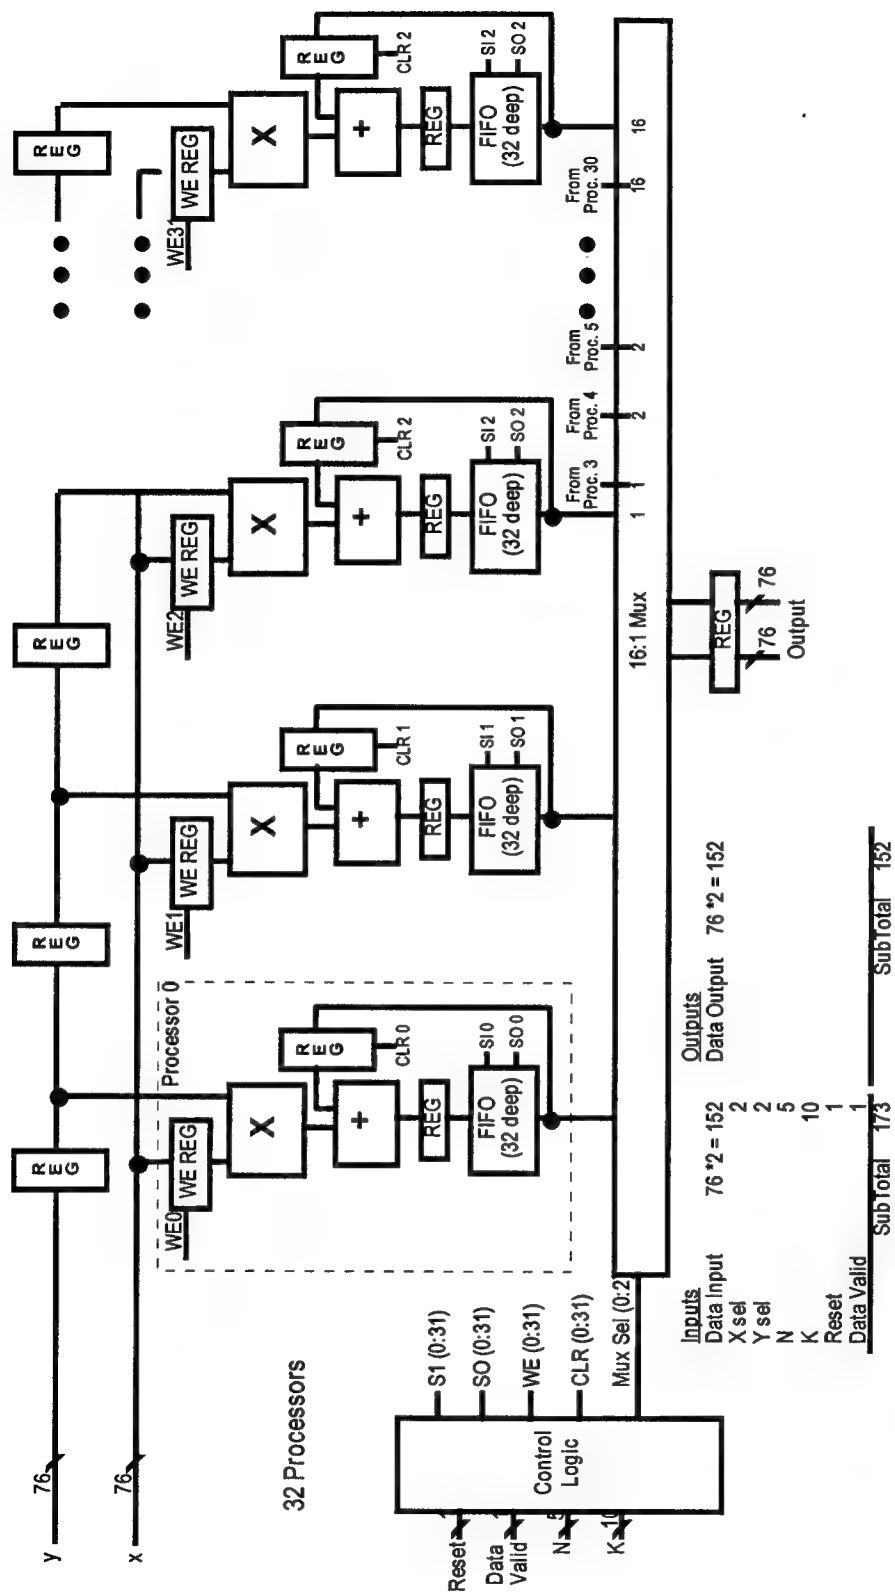


Figure 7.4.1-1 Matrix Multiplication ASIC

The scaling function is achieved through a fixed coefficient multiplication followed by RNS to binary conversion followed by binary to RNS. Threshold detection is built into the inverse square root function. The estimated gate count for this ASIC is 400K.

7.4.3 Power Dissipation For Weight Maker Module

The weight maker module will fit on a single 6U board and satisfy the IOP requirement but we must look at the power dissipation on the board. In estimating the power that an ASIC will consume, ASIC vendors will specify a number associated with their technology expressed in terms of $\mu\text{W}/\text{MHz}/\text{gate}$. A reasonable value for this number in today's technology is 0.2. This means that an ASIC that has 400K useable gates and operates at 100 MHz can be expected to dissipate about 8 Watts.

The Matrix Multiplication ASIC and the Core ASIC can both be expected to dissipate about 8 Watts per ASIC if the clock rate is 100 MHz. This is the clock rate that is assumed in Section 7.4 to derive the performance of the module. There are 20 of these ASICs on the module so that makes 160 Watts for the ASICs alone. Additional power will be consumed by the RAMs.

With conventional air cooling, the 6U board can dissipate only about 50 Watts. The most convenient solution would be to slow the clock rate of the ASICs to 25 MHz and to use four single board Weight Maker Modules to obtain the required throughput. The added latency is only 4 data matrices which is only about 2 ms. Reducing the clock rate also considerably reduces the risk in producing the custom ASICs.

7.5 BEAMFORMER MODULE

The beamformer module is required to deliver eight complex numbers at a 1 MHz rate. Each complex number output requires 128 complex multiply-adds. This amounts to an operation rate of 8.2 GOPs. This could be accomplished in COTS technology, but if the Matrix Multiply and Core ASICs are available because they are needed for the Weight Maker Module, then a Beamformer Module could be made with them. Figure 7.5-1 shows a schematic of a single board Beamformer Module.

As is the case with the Weight Maker Module, the Beamformer Module can also be constructed with 4 boards and ASICs whose clock rate is only 25 MHz.

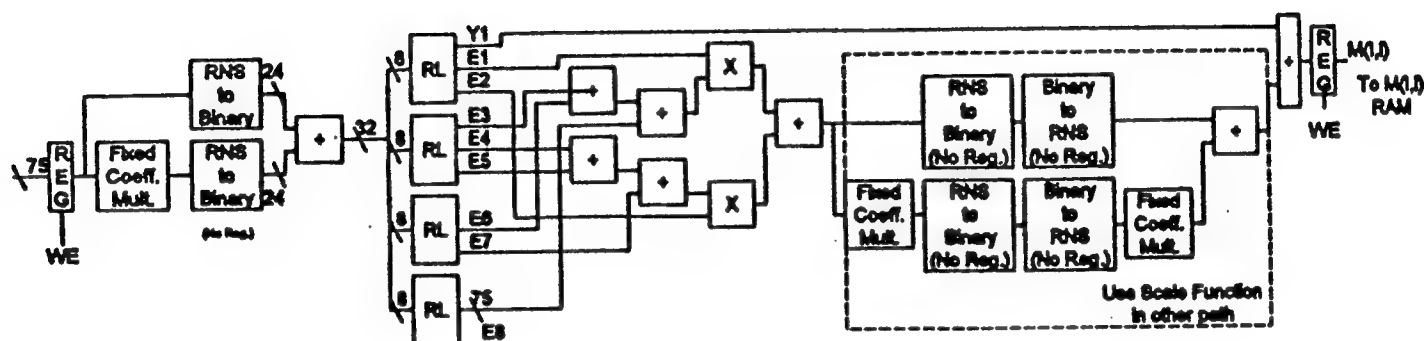
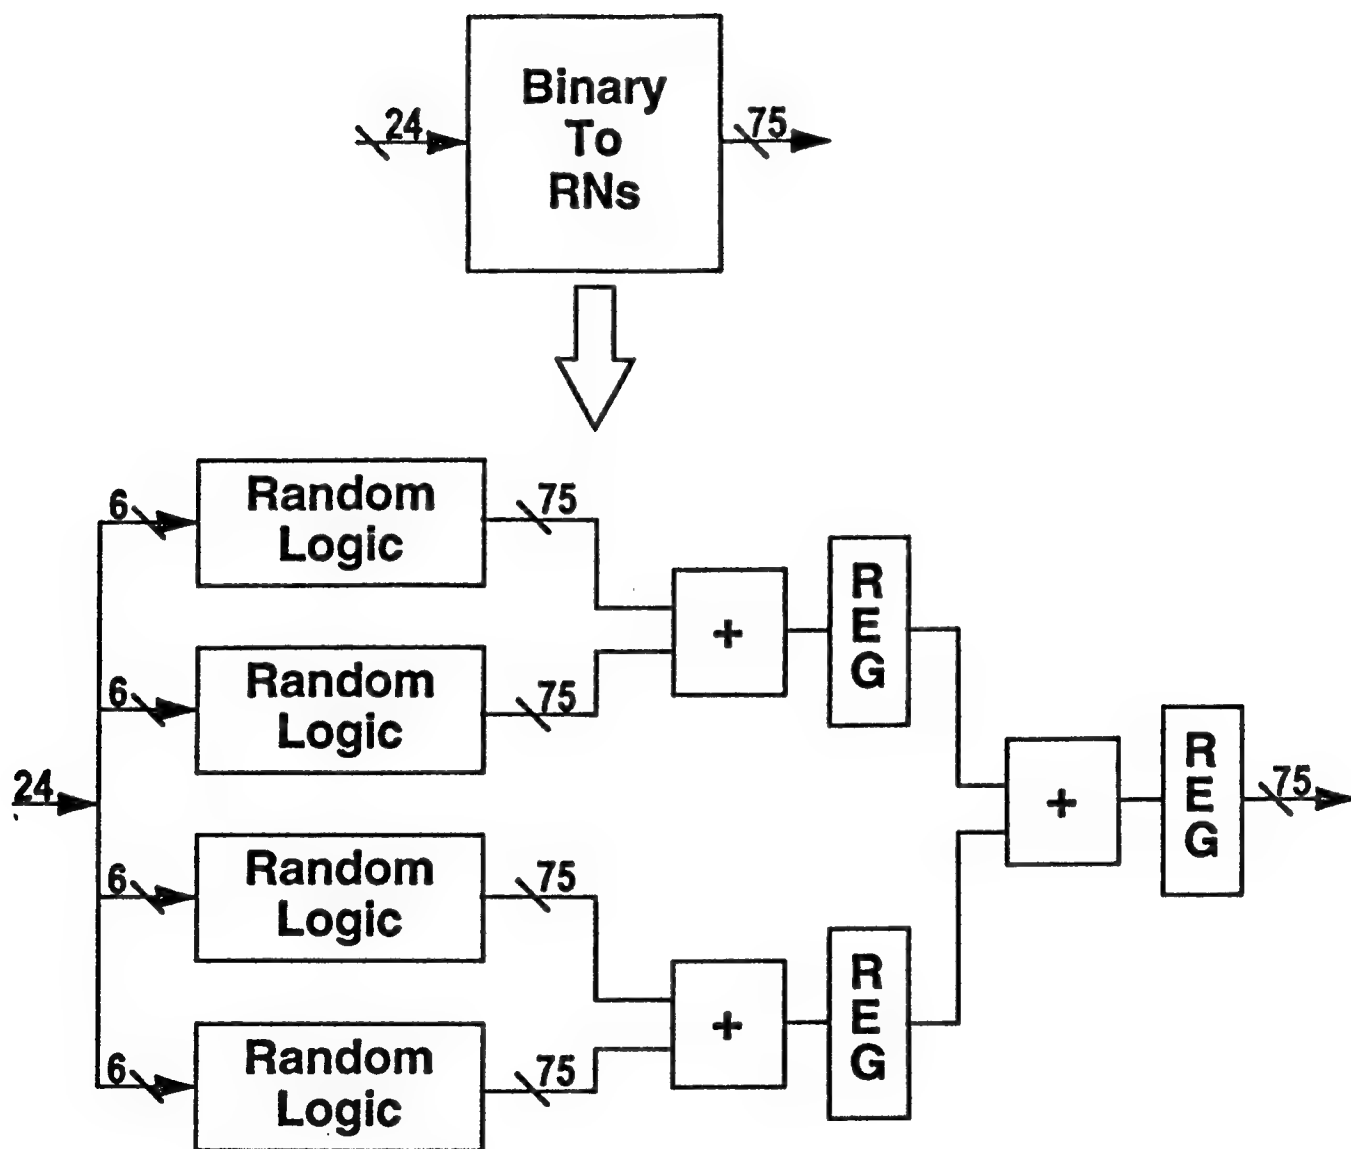


Figure 7.4.2-2 RNS Inverse Square Root



7.4.2-3 Binary-to-RNS Conversion

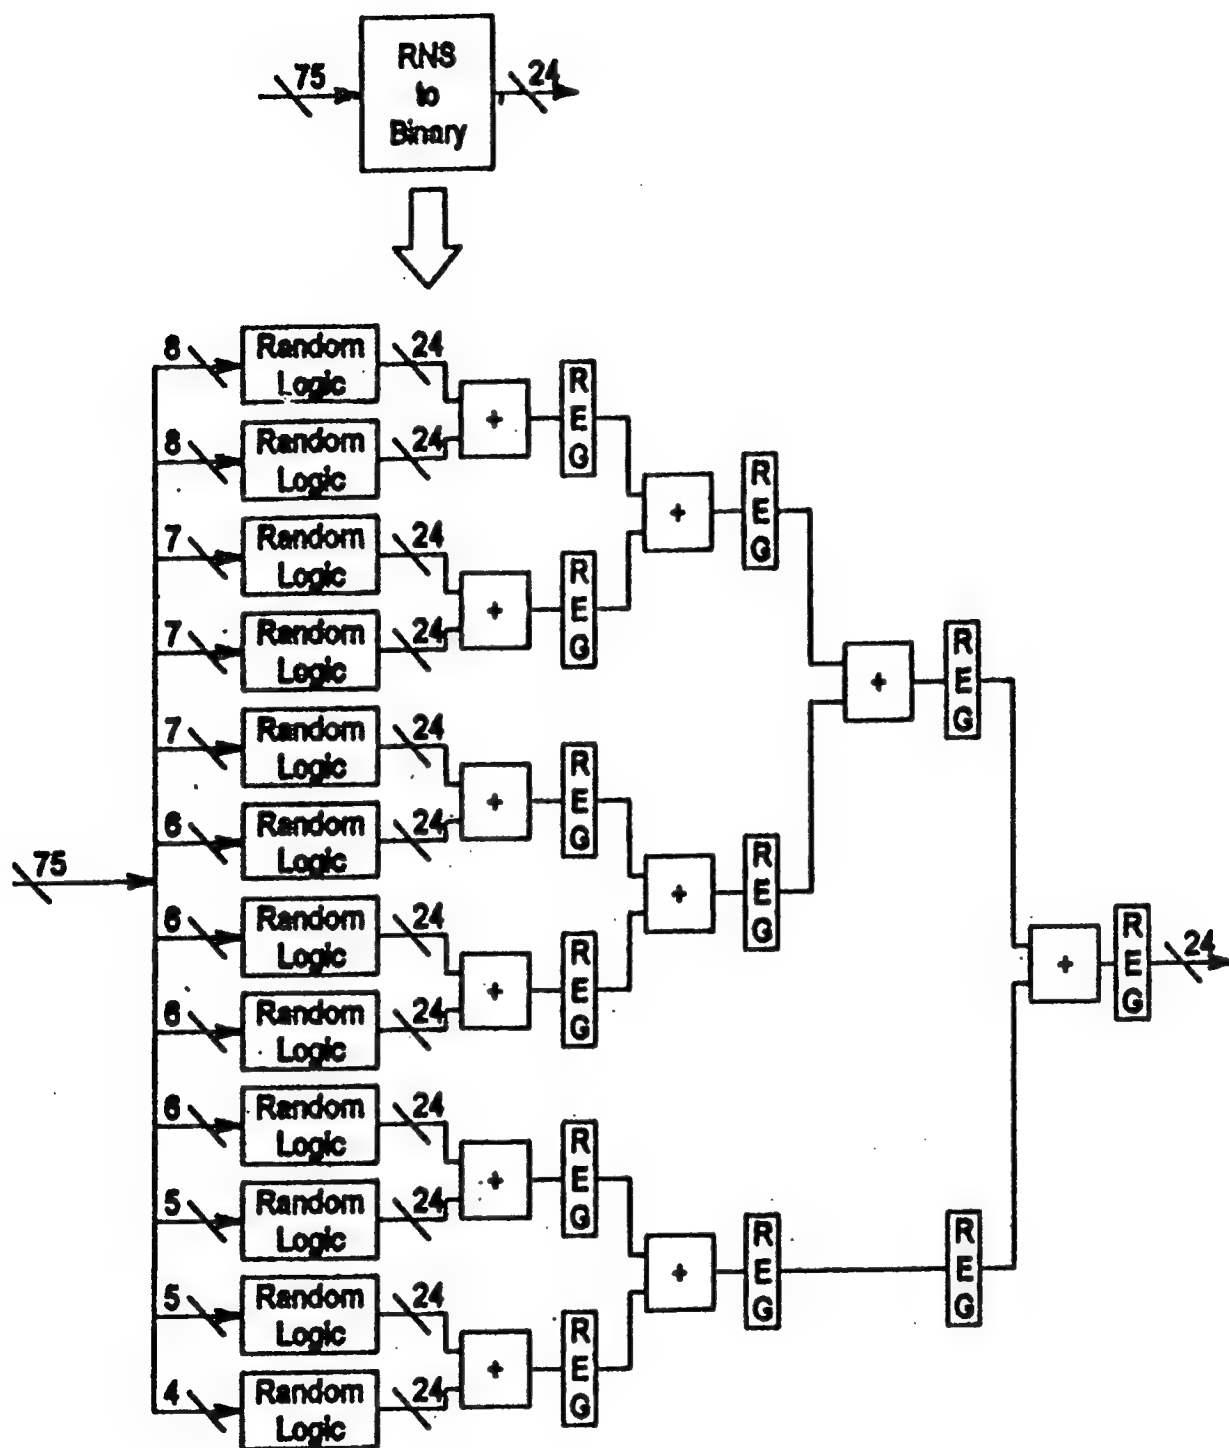


Figure 7.4.2-4 RNS-to-Binary Conversion

CRNS to QRNS




Figure 7.4.2-5(a) CRNS-to-QRNS Conversion

QRNS to CRNS

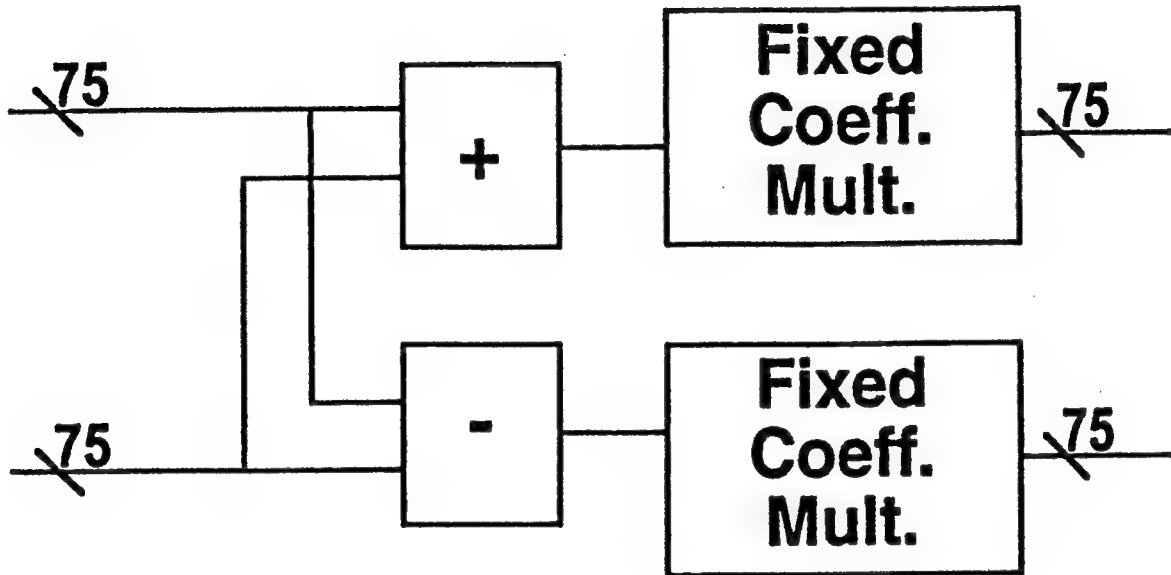


Figure 7.4.2-5(b) QRNS-to-CRNS Conversion

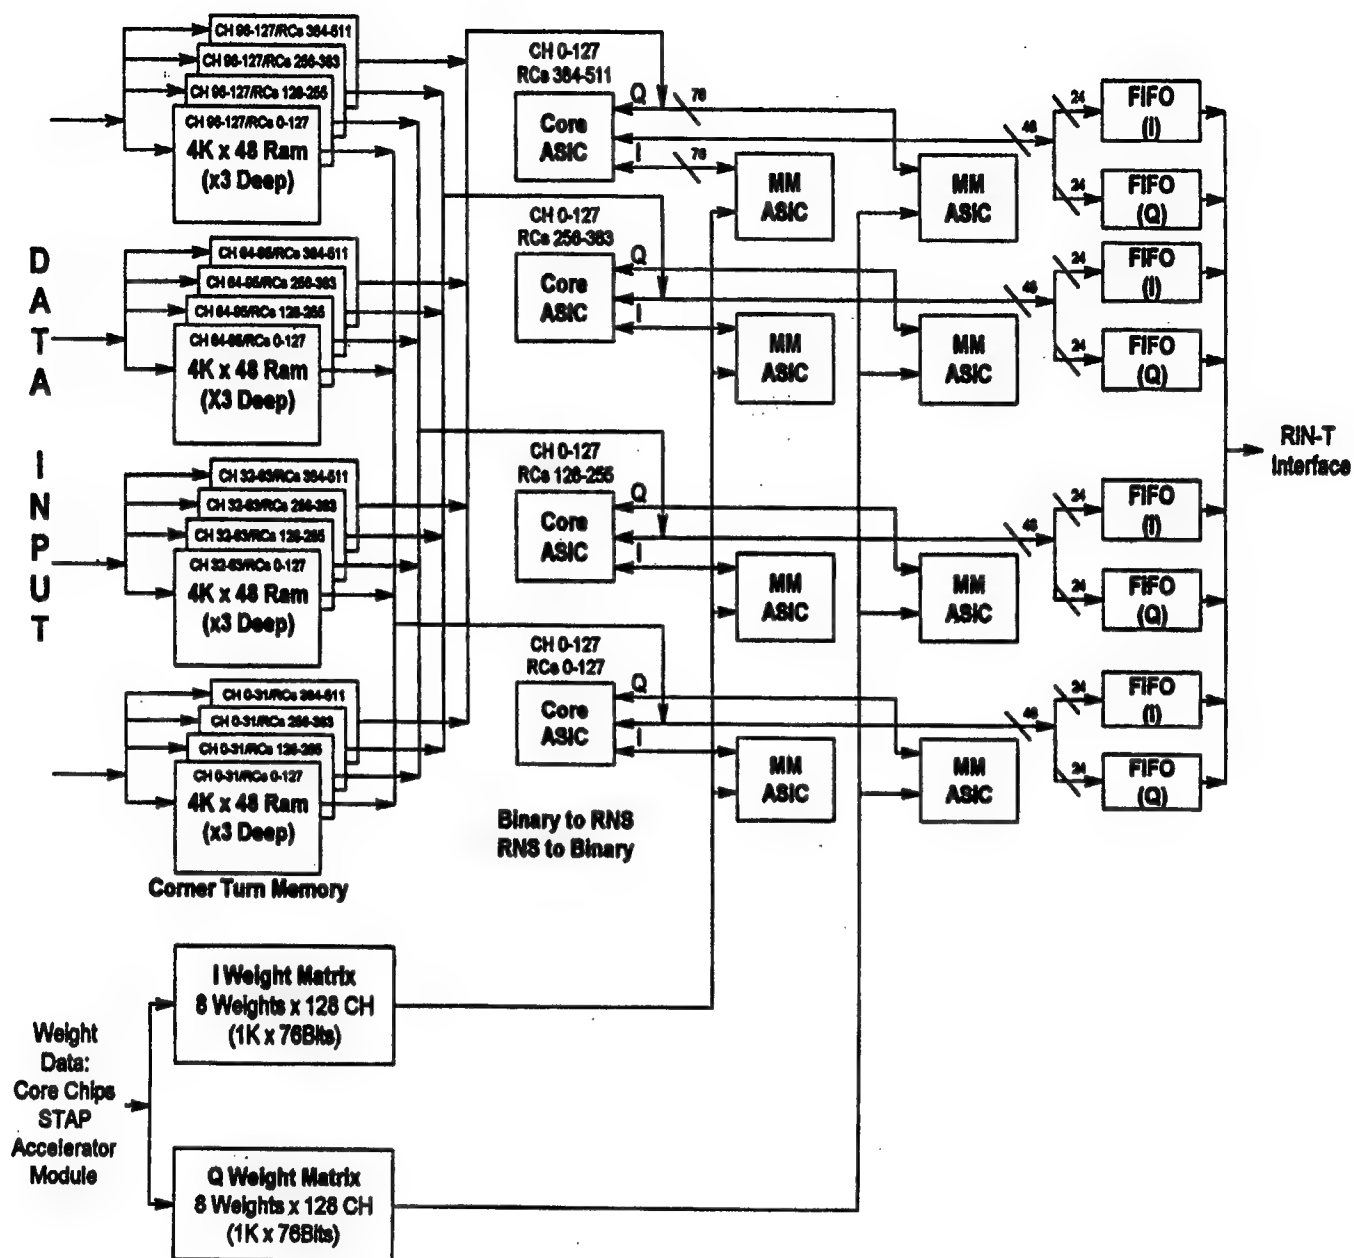


Figure 7.5-1 Beamformer Module

8. DESIGN PARAMETERS:

With the design of the IOP finalized, we are able to assess the impact of our technology insertions on our targeted STAP problem. In this section, we will measure our achieved results against those predicted by the Mountain Top program.

8.1 SYSTEM PERFORMANCE:

Tables 8.1-1 through 8.1-3 offer a comparison of our results to those predicted for the Mountain Top platform (which employs conventional COTS processing).

Since our initial examination of the Mountain Top results earlier in the IOP program, Phase II of the Mountain Top contract has been awarded to Northrop Grumman. We are thus in a position to include these more recent Phase II performance estimates in our comparison.

Table 8.1-1 illustrates the customer requirements and predicted performance of the Phase I and Phase II Mountain Top systems. (It was somewhat surprising to note that the predicted performance metrics of the Phase II Mountain Top processor, derived two years later, are actually lower than those of the Phase I processor. In discussions with the engineers performing the Mountain Top study, it was learned that the lower Phase II numbers are based upon their additional knowledge of and experience with the Mercury architecture, which led them to believe that their original Phase I predictions were probably overly optimistic.)

Table 8.1-2 summarizes our performance predictions for the IOP STAP processor. (We have chosen to err on the conservative side, and use worst case power numbers for each of our hardware modules. Thus, our final results will likely improve.)

Table 8.1-3 employs a set of performance metrics to draw comparison between our results and those of the Mountain Top program.

This comparison shows the IOP STAP processor performance to exceed that of conventional approaches to STAP by more than an order of magnitude in each category.

TABLE 8.1-1. MOUNTAIN TOP PROCESSOR ESTIMATES

	Mountain Top Phase I Requirement	Mountain Top Phase I Predicted	Mountain Top Phase II Requirement	Mountain Top Phase II Predicted
Size	20 ft3	12.6 ft3	24 ft3	14 ft3
Weight	600 lbs	300 lbs	400 lbs	400 lbs
Power	8000 W	2200 W	8000 W	6100 W
Thruput	10 Gflops (sustained)	Gflops (sustained)	20 Gflops (sustained)	21Gflops (sustained)

TABLE 8.1-2. IOP PROCESSOR ESTIMATES

	IOP Goals	IOP Predicted
Size	2 ft3	2.25 ft3
Weight	60 lbs	125 lbs
Power	800 W	760 W
Thruput	152 Gflops (sustained)	152 Gflops (sustained)

TABLE 8.1-3. PERFORMANCE COMPARISON

	Mountain Top Phase I Predicted	Mountain Top Phase II Predicted	IOP Predicted	Performance Delta IOP/Mtn Top
Gflops/ft3	1.53	1.50	67.55	45x
Gflops/lb	0.064	0.053	1.216	23x
Gflops/KW	8.77	3.44	200.0	58x

8.2 SUMMARY:

We have designed a flexible, affordable, and highly efficient COTS-based STAP processor. By incorporating key technology innovations in the areas of optical interconnects and high speed RNS processing, we have achieved several orders of magnitude improvement over the existing state-of-the-art.

Our design affords more than 150 Gigafllops of processing capability in a single 6U VME chassis (see Figure 8.2-1). We have thus opened up the potential for incorporating STAP in a wide range of applications where size, weight, and power constraints have previously rendered it unfeasible, including airborne surveillance, UAV's, and airships.

Our design is affordable enough for use in volume production.

Our design leverages and advances key innovations from DARPA's POINT (polymer waveguides) and OMNET (opto-electronics and packaging innovations) programs.

Our Mercury-RACE-based system is compatible with a large volume of existing software. This provides an upgrade path for significant airborne surveillance applications.




Figure 8.2-1 Single 6U Chassis IOP STAP Processor

The contributions of our IOP design to a reduction in system complexity as compared to that of a conventional all-electronic, all-COTS approach are illustrated in Figure 8.2-2.




Figure 8.2-2 Relative Processor Performance

MISSION OF ROME LABORATORY

Mission. The mission of Rome Laboratory is to advance the science and technologies of command, control, communications and intelligence and to transition them into systems to meet customer needs. To achieve this, Rome Lab:

- a. Conducts vigorous research, development and test programs in all applicable technologies;
- b. Transitions technology to current and future systems to improve operational capability, readiness, and supportability;
- c. Provides a full range of technical support to Air Force Material Command product centers and other Air Force organizations;
- d. Promotes transfer of technology to the private sector;
- e. Maintains leading edge technological expertise in the areas of surveillance, communications, command and control, intelligence, reliability science, electro-magnetic technology, photonics, signal processing, and computational science.

The thrust areas of technical competence include: Surveillance, Communications, Command and Control, Intelligence, Signal Processing, Computer Science and Technology, Electromagnetic Technology, Photonics and Reliability Sciences.